An Efficient Private Evaluation of a Decision Graph

Hiroki Sudo[1,2][0000−0003−2222−3371], Koji Nuida[3][0000−0001−8259−9958], and Kana Shimizu[1(✉)][0000−0001−6452−9091]

[1] Waseda University, Tokyo, Japan
hsudo108@ruri.waseda.jp, shimizu.kana@waseda.jp
[2] AIST-Waseda University CBBD-OIL, Tokyo, Japan
[3] The University of Tokyo, Tokyo, Japan
nuida@mist.i.u-tokyo.ac.jp

Abstract. A decision graph is a well-studied classifier and has been used to solve many real-world problems. We assumed a typical scenario between two parties in this study, in which one holds a decision graph and the other wants to know the class label of his/her query without disclosing the graph and query to the other. We propose a novel protocol for this scenario that can obliviously evaluate a graph that is designed by an efficient data structure called the graph level order unary degree sequence (GLOUDS). The time and communication complexities of this protocol are linear to the number of nodes in the graph and do not include any exponential factors. The experiment results revealed that the actual runtime and communication size were well concordant with theoretical complexities. Our method can process a graph with approximately 500 nodes in only 11 s on a standard laptop computer. We also compared the runtime of our method with that of previous methods and confirmed that it was one order of magnitude faster than the previous methods.

Keywords: Decision graph· Homomorphic encryption· GLOUDS.

1 Introduction

Classification is a central topic in machine learning (ML), which is aimed at training a classifier on a set of labeled samples so that the trained classifier can correctly assign one of the class labels to an input query, and has been successfully applied to various real-world problems such as credit scoring, drug discovery, and disease diagnostics [17, 20, 22]. One of the typical online services using ML is a classification service where a service provider has a trained classifier and a user can obtain classification results for his/her data. In fact, software platforms that easily achieve such scenario are already available [1–3], which enables service providers to publish the application programming interfaces (APIs) of the trained classifiers on the cloud server. Although both service providers and service users can benefit from such classification services, there are certain *privacy* concerns about the data. A natural scenario for an online classification

service is for a user to send his/her query (an input to a classifier) to a server and the server to return the classification results based on the classifier. Suppose the online service involves disease diagnostics, where the input to the classifier includes the user's private information such as health records and genetic information. The server's classifier also includes data on donors' private information because the classifier was trained on private data. Various model-inversion attacks are possible [15, 16] in this scenario; they can infer sensitive information being used for training by accessing the trained classifier. Therefore, it is necessary to conceal both the user's query and the server's classifier.

We focused on a decision graph (DG) as a classifier and tackled the problem of private evaluations of the decision graphs. A decision graph is an efficient data structure for the classification rules. It is also described as a decision diagram [7] in the logic synthesis literature and as a branching program (BP) [23] in computer science theory. Compared to complex models such as neural networks, the decision graph is easier to interpret and is therefore often preferred for problems like clinical diagnosis where the interpretation of decision-making is important. We assumed the underlying graph was a binary graph and defined a binary decision graph (BDG) as follows. BDG is a rooted directed acyclic graph (DAG) that consists of a set of nodes of in-degree ≥ 1 and the out-degree of two or zero. A node with the out-degree of zero is called a leaf and has a class label. Each internal node contains a split function that decides whether a query that reaches the node should visit a node connected to the right edge or the node connected to the left edge, depending on the corresponding attribute of the query. We assumed in our study that each split function computed whether or not the input was greater than a threshold t.

The problem setup for this study was as follows: one party (a server) has a BDG and the other party (a user) wants to obtain a classification result. The user's input is a private attribute vector, $\boldsymbol{x} = (x_0, \ldots, x_{n-1})$. The length of the vector and the ordering of the attributes are common information between the server and the user. The user only knows common information and the height of the graph (maximum path length from a root node to a leaf). After computation, the user learns the classification result (a class label); he/she does not learn anything more than what he/she already knows.

1.1 Related Works

Many studies have addressed the problem of private evaluation of classifiers [13, 21, 32]. Brickell et al. [6] and Barni et al. [4] respectively proposed methods which combine Yao's garbled circuit [36] and additively homomorphic encryption (AHE) for private evaluation of the BDG. We will present a detailed comparison of our method with those approaches in Section 4.2. Mohassel et al. [25] have also proposed a method of private evaluation of BDG; however, they assumed that a user knew all the outputs of all split functions of a server's internal nodes, which differs from our scenario.

Since the decision graph is regarded as a generalized form of a decision tree (DT), we also describe a series of studies for a private evaluation of DTs. Bost et

al. [5] proposed a secure decision tree evaluation protocol as part of their work. Their method evaluated a decision tree as a polynomial of Boolean variables using leveled fully homomorphic encryption. Although this method improved efficiency compared to other conventional methods, it still suffered from the problem of computation and communication costs. A recent work by Wu et al. [35] achieved more practical computational time and communication size. Their method was only based on AHE and performed efficiently for shallow trees; however, it did not perform well for the evaluation of deep trees because of its exponential time and communication complexity for the height of the trees. Cock et al. [10] proposed a protocol that achieved time complexity that was similar to Wu et al.'s algorithm and improved runtime by using arithmetic sharing to avoid heavy modular exponentiation. However, their protocol assumed a different problem setup where a trusted initializer participated in the protocol to generate multiplication triplets. The trusted initializer could be removed, but the additional costs of generating the multiplication triplets by the two parties was exponential to the height of the tree, which greatly deteriorated the runtime. Tai et al. [30] formulated a decision tree evaluation as a compact linear function to attain a protocol in which time complexity was only dependent on the number of internal nodes and was independent of the exponential of the tree height.

Protocols for DTs can theoretically be applied to private evaluations of BDG if the underlying graph is transformed into a tree. However, the number of nodes in the tree, that is equivalent to the graph, becomes very large. As we will discuss in Subsection 2.4, two in-coming edges to an internal node cause a copy of all the subordinates of the node on transformation, which leads to the exponential growth in total tree size.

We also noted that a BDG achieved accurate predictions while it achieved lower model complexity than DT [19, 26, 27, 29], and it even achieved considerably improved generalization [29]. A BDG with 3,000 nodes achieved the same accuracy as a DT with 22,000 nodes in the classification of a Kinect dataset in a study by Shotton et al. [29].

1.2 Our Contribution

The five main contributions of this paper are summarized below:

- We propose an efficient protocol for the oblivious evaluation of a BDG. More precisely, the protocol allows two parties, one holding a BDG T, and the other holding an attribute vector, $\boldsymbol{x}$, to determine the class label of $\boldsymbol{x}$, without revealing T and $\boldsymbol{x}$ to the other party in a semi-honest setting.
- The time and communication complexities of our protocol are linear to the number of nodes and the height of T and exclude any exponential factors.
- The DAG of the BDG in our protocol is represented by a look-up vector V, and the other party obliviously refers to V. We demonstrate how the length of V is reduced by using a succinct data structure called GLOUDS to achieve linear complexity.

- An oblivious evaluation of a split function in each internal node is conducted before graph traversal. We propose a novel protocol called eROT that enables the correct edges to obliviously be chosen during traversal.
- We implemented our protocol and tested it on BDGs of various sizes; we found that its actual runtime and communication size were concordant with the theoretical complexities. We also compared the runtime and communication size of our protocol to those in previous studies [4, 6] to confirm that our protocol was an order of magnitude faster.

The rest of the paper is organized as follows. Section 2 describes important building blocks for the proposed method and the security model that was assumed for this study. We detail our method in Section 3 and evaluate it on various datasets in Section 4. Section 5 concludes the paper.

2 Preliminary

2.1 Notation

We denote vector $\boldsymbol{v}$ as $(v_0, \ldots, v_{n-1})$ and the i-th element of $\boldsymbol{v}$ as $\boldsymbol{v}[i]$. The $\{a_0, \ldots, a_{n-1}\}$ represents a set of size n. The $\{a_i\}_{i=0}^{n-1}$ stands for $\{a_0, \ldots, a_{n-1}\}$. We define the "rotation" of a vector as: given n dimensional vector $\boldsymbol{v}$, the r-rotation of $\boldsymbol{v}$ results in vector $\hat{\boldsymbol{v}}$, each of whose elements is $\hat{\boldsymbol{v}}[(i+r)_{\mathsf{mod}\ n}] = \boldsymbol{v}[i]$. The $\langle P(x) \rangle$ returns 1 if predicate $P(x)$ is true given x, otherwise 0. The notation, $r \in_R A$, means r is a uniformly chosen random value from a set A. We define λ-bit unary representation of $x \in \{0, \ldots, \lambda - 1\}$ as a λ-bit vector that has 1 at x-th least significant bit and has 0 at the other bits, and denote it as $\mathsf{UNARY}_\lambda(x)$.

2.2 Additively Homomorphic Encryption

We used a semantically secure additively homomorphic public-key encryption scheme in our protocol and especially assumed a lifted-ElGamal cryptosystem [11] with plaintext space $\mathbb{Z}_p$ whose message in a ciphertext is located in the exponent. The public-key encryption scheme is equipped with three algorithms:

1. KeyGen**:** outputs a public/private key pair $(\mathsf{pk}, \mathsf{sk})$.
2. $\mathsf{Enc}_{\mathsf{pk}}(m)$**:** outputs a ciphertext $[\![m]\!]$, by encrypting a plaintext m, with pk.
3. $\mathsf{Dec}_{\mathsf{sk}}([\![m]\!])$**:** outputs a plaintext m, by decrypting a ciphertext $[\![m]\!]$, with sk.

$[\![m]\!]$ represents a ciphertext of a plaintext m. Likewise, $[\![\boldsymbol{v}]\!]$ represents a ciphertext vector, each of whose elements is an encryption of each element of a vector $\boldsymbol{v}$. A public key of AHE has $\oplus$, $\otimes$, $\ominus$ operations on ciphertexts described as follows. Given two plaintexts m_1, m_2, we can compute $[\![m_1 + m_2]\!] = [\![m_1]\!] \oplus [\![m_2]\!]$ by using $\oplus$ operation. We can also compute multiplication by a constant k ($[\![k \cdot m]\!] = k \otimes [\![m]\!]$). Negation on a ciphertext is represented by $\ominus [\![m]\!]$. In our setting, the user generates and holds a public/private key pair $(\mathsf{pk}, \mathsf{sk})$, and the server only receives a public key pk so that only the user can decrypt ciphertexts and the server can only conduct encryption and additively homomorphic computation.

2.3 Oblivious Transfer

Oblivious transfer (OT) is a secure two-party protocol between a sender and a chooser. A chooser in 1-out-of-N OT specifies an index $i \in \{0, \ldots, N-1\}$, and only obtains the i-th element of the sender's vector $\boldsymbol{v}$, without disclosing i to the sender. We denote the execution of OT with an index i, and a vector $\boldsymbol{v}$ by $OT_1^N(i, \boldsymbol{v})$. While there are several efficient implementations that achieve OT_1^N functionality, we use simple protocols based on additively homomorphic operation which require $O(N)$ computational cost and communication size. For the space limitation, we omit the implementation detail.

2.4 GLOUDS

The graph level order unary degree sequence (GLOUDS) [14] is the succinct data structure of a DAG, which is a query-time efficient data structure that uses the space close to information-theoretic lower bound. GLOUDS regards a DAG as an integration of a spanning tree and "non-tree" edges that are not included in the tree, and introduces the idea of "shadow nodes", which are duplicates of non-tree nodes (nodes with incoming edges > 1) to virtually treat a graph as a tree, while it avoids unnecessary copies of nodes. When we transform a DAG into an equivalent tree, it is necessary to repeat copying of a subtree rooted from a non-tree node for all the incoming edges to the node, which causes exponential growth in total tree size. Since GLOUDS generates as many shadow nodes as the number of non-tree edges, it is considered to be efficient when there are not too many non-tree edges.

More precisely, GLOUDS consists of a trit (0,1,2) sequence B of length N and an auxiliary vector H, where N is the sum of the number of nodes and the number of edges + 1. The nodes in the DAG are numbered in level order (from top to bottom and left to right) and the root is numbered 0. The nodes are visited in level order during construction of GLOUDS. When each node is visited, 0 is stored in B, and all the children of the node are stored in left-to-right order. If a child is already observed, a trit 2 is stored in B, and 1 otherwise. The root node is considered as a child of an unshown supernode, and hence 1 is stored in B as the first element (i.e., $B[0] = 1$). H memorizes numbers of shadow nodes in the order in which they appear in B as 2. For the case of the DAG in Figure 1, after storing $B[0] = 1$, $B[1] = 0$ is stored when the node "0" is visited. Since the node "1" and "2" are the children of the node "0" and they are not observed, $B[2] = 1$ and $B[3] = 1$ are stored. Similarly, $B[4] = 0$, $B[5] = 1$ and $B[6] = 1$ are stored for the visit of the node "1", and $B[7] = 0$, $B[8] = 2$ and $B[9] = 1$ are stored for the visit of the node "2". Note that $B[8] = 2$ because the node "4", which is the left child of the node "2", is already observed. 4 is recorded in $H[0] = 4$. After visiting all the nodes, B and H are described as $B = 10110110210110210000$ and $H = [4, 7]$. Either 1 or 2 in vector B corresponds to any one of the nodes in the DAG, and 0 is regarded as a delimiter between groups of siblings. Note that 0 is also considered as a parent node of a right-neighbour group of siblings; therefore, the same node appears more than once in B.

Here, we define two operations on sequence B as:

Definition 1 *Operations on trit sequence B*

$\mathsf{Rank}_c(B,p)$*: returns the number of* $c \in \{0,1,2\}$ *in the prefix* $B[0,p)$ $(0 \leq p < N)$
$\mathsf{Select}_c(B,i)$*: returns the position of the* i*-th* $c \in \{0,1,2\}$ *in* B (i *starts from* 0.)

One can move from a position p in B that stores a trit 1 or 2 (a parent node), to another position p' (x-th child of the node) by carrying out the equation below.

$$p' = \begin{cases} \mathsf{Select}_0(B, \mathsf{Rank}_1(B,p)) + x & (if\ B[p] = 1) \\ \mathsf{Select}_0(B, H[\mathsf{Rank}_2(B,p)]) + x & (if\ B[p] = 2) \end{cases} \tag{1}$$

For simplicity, we let $\mathsf{SelRan}(B,p)$ be the first term of Eq. 1. The $\mathsf{SelRan}(B,p)$ computes a position in B of the left delimiter of $B[p]$'s children. Since the siblings are stored sequentially, one can specify the x-th child by adding an offset x to $\mathsf{SelRan}(B,p)$. Figure 1 has an example of SelRan. For example, let us consider the case of $p = 2$. $B[2]$ corresponds to the node "1". The children of the node "1" are "3" and "4", and they correspond to $B[5]$ and $B[6]$. $\mathsf{SelRan}(B,2)$ returns 4 and $B[4] = 0$ is the left delimiter of them. We define a map of a position in B and a node id, such that $\mathsf{ID}(p)$ returns id of the node that corresponds to $B[p]$. For example, $\mathsf{ID}(2) = 1$ and $\mathsf{ID}(4) = 1$. Note that $\mathsf{ID}(p) = \mathsf{ID}(\mathsf{SelRan}(B,p))$. GLOUDS can be regarded as a generalization of the level order unary degree sequence (LOUDS) [18], which is a succinct data structure for ordered trees; hence, our protocol can be immediately applied to DT.

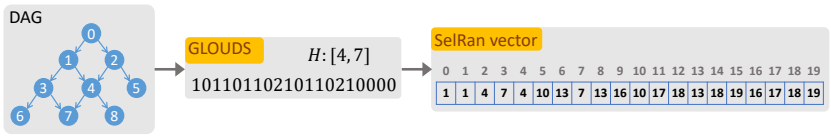

Fig. 1: Example of a BDG, corresponding GLOUDS and SelRan vector.

2.5 BDG and Efficient Design Principle of a Look-up Vector by GLOUDS

BDG consists of a rooted DAG and a set of split functions $\{\mathsf{Split}_i\}_{i=0}^{m-1}$, where m is the number of internal nodes. Given an attribute $\mathcal{X} \in \mathbb{Z}$, a split function that is assigned to an internal node performs a greater than operation: $\mathsf{Split}_i(\mathcal{X}) = \langle t_i < \mathcal{X} \rangle$ to choose either a right or left child. BDG in our protocol is mainly represented by a look-up vector $\boldsymbol{v}$ of length N (also referred to as the SelRan vector), and a vector of ciphertexts $[\![\boldsymbol{o}]\!]$ that encrypts an offset vector $\boldsymbol{o}$ of length N. $\boldsymbol{v}$ represents the DAG, and $\boldsymbol{o}$ represents outputs of all the split functions

taking a query (a set of attributes). Both $\boldsymbol{v}$ and $[\![\boldsymbol{o}]\!]$ are held by the party holding the BDG, and the other party traverses the BDG by obliviously referring to those vectors. $\boldsymbol{v}$ and $\boldsymbol{o}$ are described as:

$$\boldsymbol{v}[i] = \begin{cases} \mathsf{SelRan}(B, i) & (if\ B[i] \neq 0) \\ i & (else) \end{cases} \text{ and } \boldsymbol{o}[i] = \begin{cases} \langle \boldsymbol{\theta}[i] < \mathcal{X}_i \rangle + 1 & (if\ \boldsymbol{\tau}[i] = \mathsf{I}) \\ 0 & (else) \end{cases},$$

where $\mathcal{X}_i$ is a user's attribute for the split function that is associated with node $\mathsf{ID}(i)$, $\boldsymbol{\tau}$ is a **type vector** storing the types of each position and $\boldsymbol{\theta}$ is a **threshold vector**. $\boldsymbol{\tau}[i] = \mathsf{L}$(eaf) if node $\mathsf{ID}(i)$ is a leaf, $\boldsymbol{\tau}[i] = \mathsf{Z}$(ero) if node $\mathsf{ID}(i)$ is not a leaf and $B[i] = 0$. $\boldsymbol{\tau}[i] = \mathsf{I}$(nternal) otherwise. $\boldsymbol{\theta}[i]$ is a threshold of a split function that is associated with node $\mathsf{ID}(i)$ when $\boldsymbol{\tau}[i] = \mathsf{I}$. $\boldsymbol{\theta}[i]$ is set to empty otherwise. $\boldsymbol{v}[p]$ returns the position of the left delimiter of node $\mathsf{ID}(p)$'s children and $\boldsymbol{o}[p]$ returns the choice of a child. Therefore, one can compute the position of next node in B by:

$$p' = \boldsymbol{v}[p] + \boldsymbol{o}[p].$$

Note that the outputs of split functions include both parties' privacy; hence, the two parties need to jointly compute $[\![\boldsymbol{o}]\!]$ without revealing their private parameters. We will describe how this is accomplished in Subsection 3.4. $\boldsymbol{v}$ and $\boldsymbol{o}$ allow *self-loop* at positions $\{i \mid B[i] = 0, 0 \leq i < N\}$ by setting $\boldsymbol{v}[i] = i$ and $\boldsymbol{o}[i] = 0$. If one reaches such position i and node $\mathsf{ID}(i)$ is a leaf, one can stay on the same leaf to conceal the path length from the root toward each leaf. The self-loop at a non-leaf node can avoid incorrect traversal. The party holding BDG also prepares **label vector** $\boldsymbol{z}$. $\boldsymbol{z}[i]$ is set to a class label associated with node $\mathsf{ID}(i)$ if $\boldsymbol{\tau}[i] = \mathsf{L}$ and $B[i] = 0$. Otherwise, $\boldsymbol{z}[i]$ is a random value within the possible range of class labels. Figure 2 has an example of these data structures that represent a BDG. The nodes and edges that are colored in orange show an example of a traversal from the root node to the node 7 when $\langle t_0 < x_0 \rangle = 1$, $\langle t_2 < x_2 \rangle = 0$ and $\langle t_4 < x_4 \rangle = 0$. The corresponding elements in the table in Figure 2 are also colored in orange. The traversal starts by referring to $\boldsymbol{v}[0] = 1$ to know the next position is $\boldsymbol{v}[0] + d_0 = 3$. Similarly, one can know the next position by $\boldsymbol{v}[3] + d_2 = 8$, and visits the node 7 by $\boldsymbol{v}[8] + d_4 = 14$. Finally, one reaches the position $\boldsymbol{v}[14] + 0 = 18$ where a self-loop is allowed, and stays at the position while computing $\boldsymbol{v}[18] + 0 = 18$.

It is possible to design more space-efficient SelRan vector and auxiliary vectors, and we used it for the experiments in Section 4, however, we do not describe how we designed such vectors due to the space limitation. See the forthcoming full version of this paper for the details.

3 Method

3.1 Problem Setting

We assumed a party $\mathcal{A}$ had a private attribute vector $\boldsymbol{x}$, and a party $\mathcal{B}$ had a private BDG T, in our protocol. $\mathcal{A}$ and $\mathcal{B}$ are referred to as a user and a

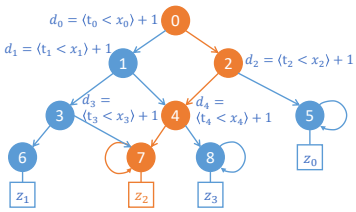

i	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
ID	0	0	1	2	1	3	4	2	4	5	3	6	7	4	7	8	5	6	7	8
B	1	0	1	1	0	1	1	0	2	1	0	1	1	0	2	1	0	0	0	0
v	1	1	4	7	4	10	13	7	13	16	10	17	18	13	18	19	16	17	18	19
z	r_0	r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}	r_{11}	r_{12}	r_{13}	r_{14}	r_{15}	z_0	z_1	z_2	z_3
τ	I	Z	I	I	Z	I	I	Z	I	L	Z	L	L	Z	L	L	L	L	L	L
θ	t_0		t_1	t_2		t_3	t_4		t_4											
o	d_0	0	d_1	d_2	0	d_3	d_4	0	d_4	0	0	0	0	0	0	0	0	0	0	0

Fig. 2: Example of a BDG (left graph) and data structures for the BDG described in Subsection 2.5 (right table). r_i is a random value.

server in previous sections. Both $\boldsymbol{x}$ and T must be concealed from the other party. $\mathcal{A}$ and $\mathcal{B}$ participate in the two-party secure BDG evaluation protocol. $\mathcal{A}$ only obtains an output of BDG $T(\boldsymbol{x})$, while he/she gains no information about $\mathcal{B}$'s private information except for $T(\boldsymbol{x})$. $\mathcal{B}$ obtains nothing. We assumed $\mathcal{A}$ and $\mathcal{B}$ shared three kinds of information: (1) length of the SelRan vector, (2) height of the BDG, and (3) length of the attribute vector. We considered a standard adversarial model in this work: a semi-honest model [8], in which even a corrupted party adheres to the specifications of a protocol.

3.2 Overview of Our Protocol

Our protocol is composed of following phases.

Comparison Phase (for constructing offset vector): $\mathcal{B}$ eventually constructs and stores the encrypted offset vector without decrypting $\mathcal{A}$'s inputs in this phase. $\mathcal{A}$ and $\mathcal{B}$ securely calculate split functions associated with nodes to achieve this purpose. $\mathcal{B}$ stores all the decisions on which branch will be selected as ciphertexts. We used a secure comparison protocol to calculate split functions.

Evaluation Phase (for computing class label on BDG): Two parties descend from the root to a leaf in the evaluation phase by recursively referring to the SelRan vector and $[\![\boldsymbol{o}]\!]$ constructed in the comparison phase. After a leaf is reached, $\mathcal{A}$ retrieves a label associated with the leaf from $\boldsymbol{z}$.

We will first describe several secure two party protocols that will be building blocks in Subsection 3.3, and then explain how to construct the comparison phase and evaluation phase in detail in Subsection 3.4. Our protocol can be seen as a sequential composition of the two protocols, Comparison Phase and Evaluation Phase. Therefore, security of our protocol is obvious if the underlying two protocols are secure.

3.3 Building Blocks

Comparison Protocol A two-party secure comparison protocol that securely computes $\langle x < y \rangle$, is required to calculate split functions in Comparison phase. We used a variant [34] of the DGK comparison protocol [12] in our implementation, which is based on additively homomorphic encryption.

While problem settings of comparison protocols vary, we assumed the following setting: $\mathcal{A}$ and $\mathcal{B}$ had a plaintext input x and y. Only $\mathcal{B}$ acquired the encrypted comparison result $[\![\langle x < y \rangle]\!]$. Since we simply used the protocol and did not modify it, we will not go into details about the specification of the comparison protocol here. When x and y are ℓ bit integers, the time complexity and communication of both $\mathcal{A}$ and $\mathcal{B}$ are $O(\ell)$ in the DGK comparison protocol.

Recursive Oblivious Transfer $\mathcal{A}$ recursively accesses $\mathcal{B}$'s SelRan vector $\boldsymbol{v}$, in the evaluation phase, i.e., $\mathcal{A}$ repeats querying an element of $\boldsymbol{v}$ and sets the next query depending on the query result. Not only queries but also intermediate results sent from $\mathcal{B}$ need to be hidden to protect private information for both parties. We used a known secure two-party protocol called recursive oblivious transfer (ROT) [28,31] for this problem.

Assuming $\mathcal{B}$ has a plaintext vector $\boldsymbol{v}$ of length N and $\mathcal{A}$ specifies a query p_0, ROT ensures that $\mathcal{A}$ obtains $\boldsymbol{v}[\boldsymbol{v}[\ldots \boldsymbol{v}[p_0]\ldots]]$ and $\mathcal{B}$ obtains nothing after an arbitrary number of iterations. ROT consists of σ steps, where σ is a common parameter between $\mathcal{A}$ and $\mathcal{B}$. Except for the initial and the last steps, rest of the steps repeat the same protocol. The initial step computes the next position starting from the initial position p_0 specified by $\mathcal{A}$. At the end of the initial step, $\mathcal{A}$ and $\mathcal{B}$ gain shares of the next position $p_1 = \boldsymbol{v}[p_0]$. The k-th step ($k = 1, \ldots, \sigma - 2$) updates the position using the shares of the k-th position p_k. i.e., $\mathcal{A}$ and $\mathcal{B}$ gain shares of the next position $p_{k+1} = \boldsymbol{v}[p_k]$, p'_{k+1} and r_{k+1} where $r_{k+1} \in_R \mathbb{Z}$ is a random value. In the last step, the final value $\boldsymbol{v}[p_\sigma]$ is not divided into shares and only $\mathcal{A}$ knows the value.

For convenience, we denote $(p'_{k+1}, r_{k+1}) \leftarrow \mathsf{ROT}(p'_k, r_k, \boldsymbol{v})$ for the k-th step of ROT, which takes shares of a query p_k (i.e., p'_k, r_k), and a vector $\boldsymbol{v}$ as inputs, and outputs p'_{k+1} to $\mathcal{A}$ and r_{k+1} to $\mathcal{B}$. The initial step can also be represented by this notation by setting 0 to r_0. The time complexity and communication size of one round in ROT is $O(N)$ on both $\mathcal{A}$'s and $\mathcal{B}$'s sides due to the cost of OT.

eROT The goal of eROT is recursive references to the offset vector $\boldsymbol{o}$ when it is encrypted. Specifying a query p_0, $\mathcal{A}$ obtains $\boldsymbol{o}[\boldsymbol{o}[\ldots \boldsymbol{o}[p_0]\ldots]]$ after an arbitrary number of iterations, and $\mathcal{B}$ obtains nothing. To achieve this goal, we assumed $\mathcal{B}$ had an $N \times \lambda$ ciphertext matrix $\boldsymbol{\Omega}$, instead of $[\![\boldsymbol{o}]\!]$. Each row $\boldsymbol{\Omega}[i]$ is meant to represent $\boldsymbol{o}[i]$. More concretely, $\boldsymbol{\Omega}[i]$ is a vector, each of whose elements is an encrypted bit of $\mathsf{UNARY}_\lambda(\boldsymbol{o}[i])$. For example, $\boldsymbol{o}[i] = 0$ is represented by $\boldsymbol{\Omega}[i][0] = [\![1]\!]$, $\boldsymbol{\Omega}[i][1] = [\![0]\!]$ and $\boldsymbol{\Omega}[i][2] = [\![0]\!]$ when $\lambda = 3$.

Initial (0-th) step:

1. $\mathcal{B}$ generates a random value $r_1 \in_R \mathbb{Z}$, and prepares a vector $\boldsymbol{u}'$ whose i-th element is masked by r_1: $\boldsymbol{u}'[i] = \bigoplus_{j=0}^{\lambda-1}(\boldsymbol{\Omega}[i][j] \otimes (j + r_1)_{\mathsf{mod}\ N})$ (namely, $\boldsymbol{u}[i] = [\![\boldsymbol{o}[i]]\!], \boldsymbol{u}'[i] = [\![(\boldsymbol{o}[i] + r_1)_{\mathsf{mod}\ N}]\!]$.) $\mathcal{B}$ stores r_1.
2. $\mathcal{A}$ (chooser) and $\mathcal{B}$ (sender) engage in $OT_1^N(p_0, \boldsymbol{u}')$. $\mathcal{A}$ obtains $(p_1 + r_1)_{\mathsf{mod}\ N}$ decrypting $\boldsymbol{u}'[p_0] = [\![(p_1 + r_1)_{\mathsf{mod}\ N}]\!]$.

k-th ($k = 1, \ldots, \sigma - 2$) step:
$\mathcal{A}$ holds $p'_k = (p_k + r_k)_{\mathsf{mod}\ N}$ and $\mathcal{B}$ holds r_k.

1. $\mathcal{B}$ generates a random value $r_{k+1} \in_R \mathbb{Z}$. Then, $\mathcal{B}$ prepares a vector $\boldsymbol{u}'$ whose i-th element is masked by r_{k+1}: $\boldsymbol{u}'[i] = \bigoplus_{j=0}^{\lambda-1}(\boldsymbol{\Omega}[i][j] \otimes (j + r_{k+1})_{\mathsf{mod}\ N})$. $\mathcal{B}$ stores r_{k+1}.
2. $\mathcal{B}$ rotates $\boldsymbol{u}'$ by r_k elements to obtain $\hat{\boldsymbol{u}}'$.
3. $\mathcal{A}$ (chooser) and $\mathcal{B}$ (sender) engage in $OT_1^N(p'_k, \hat{\boldsymbol{u}}')$, and $\mathcal{A}$ obtains $(p_{k+1} + r_{k+1})_{\mathsf{mod}\ N}$ decrypting $\boldsymbol{u}'[p_k] = [\![(p_{k+1} + r_{k+1})_{\mathsf{mod}\ N}]\!]$.

Last step:
$\mathcal{B}$ does not mask $\boldsymbol{u}[i](= \bigoplus_{j=0}^{\lambda-1}(\boldsymbol{\Omega}[i][j] \otimes j_{\mathsf{mod}\ N}))$ in the last step to send a true value to $\mathcal{A}$.

1. $\mathcal{B}$ rotates $\boldsymbol{u}$ by $r_{\sigma-1}$ elements to obtain $\hat{\boldsymbol{u}}$.
2. $\mathcal{A}$ (chooser) and $\mathcal{B}$ (sender) engage in $OT_1^N(p'_{\sigma-1}, \hat{\boldsymbol{u}})$, and $\mathcal{A}$ obtains $\boldsymbol{u}[p_{\sigma-1}] = [\![p_\sigma]\!]$. $\mathcal{A}$ obtains $\boldsymbol{o}[\boldsymbol{o}[\ldots \boldsymbol{o}[p_0] \ldots]] = p_\sigma$ by decrypting $[\![p_\sigma]\!]$.

For convenience, we denote $(p'_{k+1}, r_{k+1}) \leftarrow \mathsf{eROT}(p'_k, r_k, \boldsymbol{\Omega})$ for the k-th step of eROT, which takes shares of a query p_k (i.e., p'_k, r_k) and a matrix $\boldsymbol{\Omega}$ as inputs, and outputs p'_{k+1} to $\mathcal{A}$ and r_{k+1} to $\mathcal{B}$. The initial step can also be represented by this notation setting from 0 to r_0. Since the major part of the time complexity is the inner product and OT, the time complexity on $\mathcal{B}$'s side in one round is $O(N\lambda)$. The time complexity on $\mathcal{A}$'s side is $O(N)$ per iteration. The communication size per iteration is $O(N)$ due to the communication size for OT.

We state that the following security theorem is established for eROT.

Theorem 1. eROT *correctly outputs* $\boldsymbol{o}[\boldsymbol{o}[\ldots \boldsymbol{o}[p_0] \ldots]]$ *and is secure in the semi-honest setting.*

Proof. **Correctness:** Each row of $\boldsymbol{\Omega}$ is an unary representation of a value, and hence conducting $\boldsymbol{u}'[i] = \bigoplus_{j=0}^{\lambda-1}(\boldsymbol{\Omega}[i][j] \otimes (j+r)_{\mathsf{mod}\ N})$ correctly yields an encryption of $(p+r)_{\mathsf{mod}\ N}$, where p is the value stored at $\boldsymbol{\Omega}[i]$. Therefore, by performing the initial step of eROT, the two parties can obtain shares $(p_1 + r_1)_{\mathsf{mod}\ N}$ and r_1 of the true position p_1. In the k-th step, $\mathcal{A}$'s input $p'_k = (p_k + r_k)_{\mathsf{mod}\ N}$ to OT is a share of the true position p_k, and the two parties can obtain a correct element by rotating $\boldsymbol{u}'$ by r_k before conducting OT. After decrypting the encrypted value obtained by OT, $\mathcal{A}$ knows the share of the next position p_{k+1}. In the last step, $\mathcal{B}$ does not mask $\boldsymbol{u}'$ and, therefore by induction it holds that the protocol correctly outputs $\boldsymbol{o}[\boldsymbol{o}[\ldots \boldsymbol{o}[p_0] \ldots]]$ to $\mathcal{A}$.
Security: All the messages are exchanged by OT. Considering that secure OT is used, it is guaranteed that no information of $\mathcal{A}$ is leaked to $\mathcal{B}$. Security against a semi-honest user is established by secret sharing. Shares of intermediate results received by $\mathcal{A}$ are indistinguishable from uniformly distributed random values due to the property of modular addition. Thus, a semi-honest user cannot acquire any information from intermediate results. □

Algorithm 1 Detailed description of comparison phase

- Public inputs: length of GLOUDS N; height d; length of attribute vector n
- Private input of $\mathcal{B}$: threshold vector $\boldsymbol{\theta}$; type vector $\boldsymbol{\tau}$
- Private input of $\mathcal{A}$: attribute vector $\boldsymbol{x}$

Step (1): $\mathcal{A}$ and $\mathcal{B}$ conduct comparison protocol coorperatively and $\mathcal{B}$ obtains $[\![\langle \boldsymbol{\theta}[j] < \mathcal{X}_j \rangle]\!]$. $\mathcal{X}_j$ is an element of attribute vector corresponding to the position j. $\mathcal{B}$ constructs a flag matrix $\boldsymbol{F}$.

for $j \in \{0, \ldots, N-1\}$ **do**
  if $\boldsymbol{\tau}[j] = \mathsf{I}$ **then**
    $\boldsymbol{F}[j][0] \leftarrow [\![1]\!]$, $\boldsymbol{F}[j][1] \leftarrow [\![\langle \boldsymbol{\theta}[j] < \mathcal{X}_j \rangle]\!]$, $\boldsymbol{F}[j][2] \leftarrow [\![0]\!]$

Step (2): $\mathcal{B}$ constructs $\boldsymbol{W}$ from $\boldsymbol{F}$

for $j \in \{0, \ldots, N-1\}, k \in \{1, 2\}$ **do**
  if $\boldsymbol{\tau}[j] = \mathsf{I}$ **then**
    $\boldsymbol{W}[j][k] \leftarrow \boldsymbol{F}[j][k-1] \oplus (\ominus \boldsymbol{F}[j][k])$ ▷ $[\![\mathsf{UNARY}_2(\langle \boldsymbol{\theta}[j] < \mathcal{X}_j \rangle)]\!]$

Step (3): $\mathcal{B}$ constructs an encrypted offset matrix $\boldsymbol{\Omega}$ based on $\boldsymbol{W}$.

for $j \in \{0, \ldots, N-1\}$ **do**
  if $\boldsymbol{\tau}[j] = \mathsf{Z}$ or $\boldsymbol{\tau}[j] = \mathsf{L}$ **then**
    $\boldsymbol{\Omega}[j][0] \leftarrow [\![1]\!]$, $\boldsymbol{\Omega}[j][1] \leftarrow [\![0]\!]$, $\boldsymbol{\Omega}[j][2] \leftarrow [\![0]\!]$ ▷ $[\![\mathsf{UNARY}_3(0)]\!]$
  else if $\boldsymbol{\tau}[j] = \mathsf{I}$ **then**
    $\boldsymbol{\Omega}[j][0] \leftarrow [\![0]\!]$, $\boldsymbol{\Omega}[j][1] \leftarrow \boldsymbol{W}[j][1]$, $\boldsymbol{\Omega}[j][2] \leftarrow \boldsymbol{W}[j][2]$
    ▷ $[\![\mathsf{UNARY}_3(\langle \boldsymbol{\theta}[j] < \mathcal{X}_j \rangle + 1)]\!]$

3.4 Secure BDG Evaluation Using GLOUDS and AHE

Comparison Phase $\mathcal{A}$ and $\mathcal{B}$ construct an encrypted matrix $\boldsymbol{\Omega}$, which corresponds to offset vector $\boldsymbol{o}$. The comparison phase ensures that no information from $\mathcal{B}$ or $\mathcal{A}$ will be disclosed, other than the number of comparisons. The following describes how we constructed $\boldsymbol{\Omega}$. The detailed algorithm is provided in Algorithm 1.

Construction of $\boldsymbol{F}$: Each split function is associated with one of positions $\{j \mid \boldsymbol{\tau}[j] = \mathsf{I} \wedge 0 \leq j < N\}$. $\mathcal{A}$ and $\mathcal{B}$ conduct a secure comparison protocol in Step (1) of Algorithm 1 to securely compute all the comparison results between attributes and thresholds. $\mathcal{B}$ finally constructs a flag matrix $\boldsymbol{F}$. We do not need to compute $\boldsymbol{F}[j]$ if $\boldsymbol{\tau}[j] \neq \mathsf{I}$.

Construction of $\boldsymbol{W}$: $\mathcal{B}$ constructs a matrix $\boldsymbol{W}$, each of whose rows is an encrypted 2-bit unary vector that represents an output of a split function.

Construction of $\boldsymbol{\Omega}$: $\mathcal{B}$ constructs $\boldsymbol{\Omega}$ in Step (3) based on $\boldsymbol{W}$ and $\boldsymbol{\tau}$. To make an encryption of $\mathsf{UNARY}_3(\langle \boldsymbol{\theta}[j] < \mathcal{X}_j \rangle + 1)$, we set $\boldsymbol{\Omega}[j][0] = [\![0]\!]$ if $\boldsymbol{\tau}[j] = \mathsf{I}$. If $\boldsymbol{\tau}[j] \neq \mathsf{I}$, $\boldsymbol{\Omega}[j]$ stores $\mathsf{UNARY}_3(0)$. Note that the lengths of rows of $\boldsymbol{\Omega}$ can be reduced to 1 when $\tau[j] \neq \mathsf{I}$ (because the offset is 0). This is because $\mathcal{B}$ knows the offsets that do not rely on any user information and can minimize the bit length. To use the reduced form of the offset matrix, we modify the inner product in

each step of eROT to $\bigoplus_{j=0}^{u_j}(\boldsymbol{\Omega}[i][j]\cdot(j+r)_{\mathsf{mod}\ N})$, where $u_j \in \{1,3\}$ is the length of the row. As a result, we can also reduce the time complexity to $O(N)$, where N is the length of GLOUDS.

Also note that the calculation of an offset can be omitted when the node is a shadow node. A new type of position for shadow nodes should be defined to do that to distinguish them from other internal nodes.

We state that the following security theorem is established for Algorithm 1.

Theorem 2. *Algorithm 1 correctly outputs $\boldsymbol{\Omega}$ and is secure in the semi-honest setting.*

Proof. **Correctness:** When an attribute $\mathcal{X}_j$ is less than an threshold $\boldsymbol{\theta}[j]$, $(\boldsymbol{W}[j][1], \boldsymbol{W}[j][2])$ becomes $(\llbracket 1 \rrbracket, \llbracket 0 \rrbracket)$, otherwise $(\llbracket 0 \rrbracket, \llbracket 1 \rrbracket)$ assuming the correctness of the underlying secure comparison protocol. Therefore, all the rows of $\boldsymbol{\Omega}$ satisfy the condition that they represent offsets in encrypted unary vectors.
Security: We have assumed that the underlying secure comparison protocol is secure in the semi-honest setting. Since the procedures after the secure comparison protocol only require server side operations on ciphertext, the security of the comparison phase is guaranteed by the security of the secure comparison protocol. □

Evaluation Phase This section describes the evaluation phase in which the participants securely descend a BDG using ROT and eROT.

We need to recursively refer to $\boldsymbol{v}$ and $\boldsymbol{\Omega}$ by starting from an initial position $p_0 = 0$ to move from the root to the leaf. The next position p_{i+1}, given a starting position p_i on GLOUDS, which corresponds to the child, is calculated by adding the p_i-th elements of a SelRan vector $\boldsymbol{v}$, and an encrypted offset matrix $\boldsymbol{\Omega}$. The next iteration will be executed after the next position p_{i+1} is set.

The private information of both parties must simultaneously be protected. There are two main security requirements: (1) $\mathcal{B}$ should not know the positions specified by $\mathcal{A}$ or the results of the protocol, and (2) $\boldsymbol{v}$ and $\boldsymbol{\Omega}$ held by $\mathcal{B}$ should be concealed from $\mathcal{A}$. We used ROT and eROT to recursively refer to $\boldsymbol{v}$ and $\boldsymbol{\Omega}$ concealing private information. Algorithm 2 describes the details of the evaluation phase satisfies the previously explained functionality and security requirements.

First, $\mathcal{A}$ and $\mathcal{B}$ start initialization in Step (1). $\mathcal{B}$ sets $r_1 = r_2 = r' = 0$. The r_1, r_2 store random values used in the previous iterations of ROT and eROT, and r' is sum of r_1 and r_2 modulo N. $\mathcal{A}$ sets the initial position $p'_0 = 0$. The two parties engage in ROT and eROT in Step (2) to update the position on GLOUDS by recursively referring to $\boldsymbol{v}$ and $\boldsymbol{\Omega}$. The (β'_{k+1}, r_1) and (ω'_{k+1}, r_2) correspond to random shares of $\boldsymbol{v}[p_{k+1}]$ and an offset $\boldsymbol{o}[p_{k+1}]$, which can be recovered by using $\mathcal{B}$'s random values r_1, r_2, and rotating $\boldsymbol{v}$ and $\boldsymbol{\Omega}$ (without knowing these values due to the security of OT). Since the position of the x-th child is determined by $\mathsf{SelRan}(B, p) + x$, the next query is $p'_{k+1} = (\beta'_{k+1} + \omega'_{k+1})_{\mathsf{mod}\ N}$. This iteration is conducted d times regardless of the depth of a leaf, which should be reached. It should be noted that a position is fixed once it is reached at a position in B with

Algorithm 2 Detailed description of evaluation phase

– Public input: length of GLOUDS N; height d
– Private input of $\mathcal{B}$: SelRan vector $\boldsymbol{v}$, encrypted offset matrix $\boldsymbol{\Omega}$, label vector $\boldsymbol{z}$

Step (1): Initialization
$\mathcal{B}$ conducts: $r_1 \leftarrow 0$, $r_2 \leftarrow 0$, $r' \leftarrow r_1 + r_2$
$\mathcal{A}$ conducts: $p'_0 \leftarrow 0$

Step (2): Update the position in GLOUDS by iterating ROT, eROT.
for $k = 0$ **to** $d - 1$ **do**
$\mathcal{A}$ and $\mathcal{B}$ engage in ROT and eROT.
$(\beta'_{k+1}, r_1) \leftarrow \mathsf{ROT}(p'_k, r', \boldsymbol{v})$, $(\omega'_{k+1}, r_2) \leftarrow \mathsf{eROT}(p'_k, r', \boldsymbol{\Omega})$
$\mathcal{B}$ conducts: $r' \leftarrow (r_1 + r_2)_{\mathsf{mod}\ N}$
$\mathcal{A}$ conducts: $p'_{k+1} \leftarrow (\beta'_{k+1} + \omega'_{k+1})_{\mathsf{mod}\ N}$

Step (3): Get the output of BDG $T(\boldsymbol{x})$ from $\boldsymbol{z}$ using $OT_1^N(p'_d, \boldsymbol{z})$.

trit 0, which ensures that the last position is in the position that corresponds to the leaf due to the definition of $\boldsymbol{v}$ and $\boldsymbol{\Omega}$. Finally, $\mathcal{A}$ obtains the output of the BDG $T(\boldsymbol{x})$ from $\boldsymbol{z}$ using $OT_1^N(p'_d, \boldsymbol{z})$ in Step (3).

We state that the following security theorem is established for Algorithm 2.

Theorem 3. *Algorithm 2 correctly outputs $T(\boldsymbol{x})$ and is secure in the semi-honest setting.*

Proof. **Correctness:** Due to the way the look-up vector $\boldsymbol{v}$ and $\boldsymbol{\Omega}$ are constructed, it is obvious that the evaluation phase can correctly compute $\boldsymbol{v}[p_k] + \boldsymbol{o}[p_k]$ in k-th step, if $\boldsymbol{v}$ and $\boldsymbol{\Omega}$ are not randomized. $\boldsymbol{v}$ is randomized by r_1 and $\boldsymbol{\Omega}$ (i.e., $\boldsymbol{o}$) is randomized by r_2. Since following equation is established by considering the property of modular addition,

$$\begin{aligned}\{p'_{k+1} - r'\}_{\mathsf{mod}\ N} &= \{(\beta'_{k+1} - r_1) + (\omega'_{k+1} - r_2)\}_{\mathsf{mod}\ N} \\ &= \boldsymbol{v}[p_k] + \boldsymbol{o}[p_k] = p_{k+1},\end{aligned}$$

$p'_{k+1} = (\beta'_{k+1} + \omega'_{k+1})_{\mathsf{mod}\ N}$ and $r' = (r_1 + r_2)_{\mathsf{mod}\ N}$ are the shares of p_{k+1}, hence one can obtain shares of $\boldsymbol{v}[p_{k+1}]$ and $\boldsymbol{o}[p_{k+1}]$ by conducting ROT and eROT with the same arguments p'_{k+1} and r'. The label corresponding to the leaf is obtained by OT in Step (3). Therefore, by induction it holds that the evaluation phase correctly outputs $T(\boldsymbol{x})$.
Security: For the space limitation, we will only sketch out a proof. In Algorithm 2, all the messages are exchanged through OT in the subroutines ROT and eROT. Considering secure OT is used, it is guaranteed that no information of $\mathcal{A}$ is leaked to $\mathcal{B}$. All the messages sent from $\mathcal{B}$ are random share of $\mathcal{B}$'s information. Therefore, it is guaranteed that no information of $\mathcal{B}$ is leaked to $\mathcal{A}$ except for the final output $T(\boldsymbol{x})$. □

3.5 Complexity

This subsection discusses the asymptotic time complexity and communication complexity of our protocol. The majority of computational and communication costs in the comparison phase are due to the comparison protocol and the construction of an encrypted offset matrix. The time complexity of the comparison protocol on the whole is $O(\ell m)$ on $\mathcal{B}$'s side. It is $O(\ell(n+m))$ on $\mathcal{A}$'s side by considering the encryption of an attribute vector and decryption of intermediate results. The construction of an encrypted offset matrix requires $O(N)$ computational cost as the computational cost is linear to the sum of the lengths of its rows. Therefore, the total time complexity of the comparison phase is $O(\ell m+N)$ on $\mathcal{B}$'s side and $O(\ell(n+m))$ on $\mathcal{A}$'s side. Since an encrypted offset matrix is constructed offline, all of the communication cost is required by the secure comparison protocol. The communication cost for $\mathcal{B}$ is $O(\ell m)$ and that for $\mathcal{A}$ is $O(\ell n+m)$ in the comparison phase. The time complexities of both $\mathcal{A}$ and $\mathcal{B}$ are $O(dN)$ for the evaluation phase. This is because ROT and eROT require $O(dN)$ computational cost. The communication cost is also $O(dN)$. We have summarized the time and communication complexity in Table 1.

Table 1: Time complexity and Communication of each phase of our method. d is the height of a DAG, ℓ is the bit length of $\mathcal{A}$'s and $\mathcal{B}$'s inputs, m is the number of split functions, n is the number of nodes in a DAG and N is the length of B.

Phase	Time	Communication	# rounds
Comparison($\mathcal{A}$)	$O(\ell(n+m))$	$O(\ell n)$	2
Comparison($\mathcal{B}$)	$O(\ell m+N)$	$O(\ell m)$	
Evaluation($\mathcal{A}$)	$O(dN)$	$O(dN)$	$d+1$
Evaluation($\mathcal{B}$)	$O(dN)$	$O(dN)$	

4 Experiments

We evaluated the efficiency of our protocol with experiments under various settings. We implemented our protocol, which is secure against the semi-honest model using the C++ library of elliptic curve ElGamal encryption [24]. We used secp256k1 for the security parameters of lifted-ElGamal, which is secure at the 128-bit security level [9]. We used a standard desktop PC with a Xeon 3.40-GHz processor for a party $\mathcal{B}$ (a server) and a standard desktop PC with a Xeon 2.40-GHz processor for a party $\mathcal{A}$ (a user) (1 thread each). Both the server and the user were in the same local area network (LAN) in our experiments.

4.1 Experiment on Simulated Dataset

First, we will present the results obtained from experiments on the simulated dataset. The dataset was composed of pairs of BDGs and attribute vectors. We

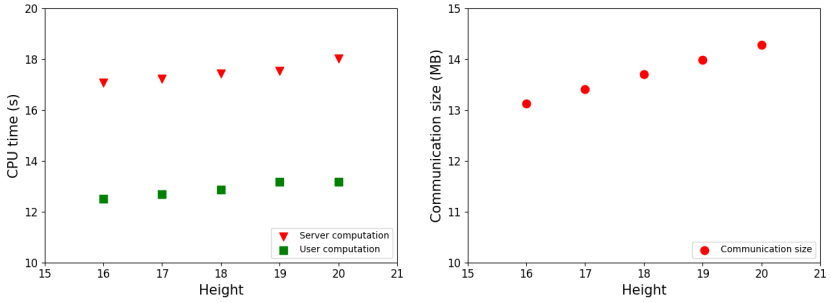

Fig. 3: CPU time (s) of server and user, and communication size (MB) on simulated datasets. We varied d from 16 to 20 while fixing other parameters ($m = 557, n = 95$, and $N = 1110$.)

varied the height d of the BDGs from 16 to 20 one by one, while fixing the number of nodes to 1110 (the number of comparisons m was 557.) and the lengths of attribute vectors n to 395. These parameters (except for d) were taken from Brickell et al. [6] to enable performance to be later compared in Section 4.2. Figures 3 plots the CPU time and communication size of the server and the user in our protocol. We observed that even when d was 20, our protocol finished within a practical timeframe and communication size (27 s and 14 MB). We also confirmed that the CPU time and communication size of both the server and the user were linear to d, which is consistent with theoretical complexity. We also confirmed that runtime overhead caused by network latency was not too large. When the round trip time (RTT) was 45 ms, which is regional RTT within North America [33], the increase in runtime was only 6 s.

4.2 Comparison to Conventional Methods

Brickell et al. [6] proposed a $O(n + \ell N + d)$ time BP (BDG) evaluation method based on AHE and Garbled Circuit. [6] reported the performance of their protocol on a BDG (1107 nodes and 395 attributes) as 302 s in CPU time and 25 MB in communication size. Since the exact topology of the BDG they used was unshown, we conducted an experiment on BDGs that have the same number of nodes and attributes with various heights for fair comparison. The results revealed that our method maintained 11 times better performance in runtime (26 s) than that of [6] and required 1.8 times less communication size (14 MB) even when the DAG is as high as $d = 20$. Barni et al. proposed a privacy-preserving evaluation method of LBP which is a generalization of BP. The time complexity of this method is $O(n+m\ell'+d)$ where $\ell'(> \ell)$ is bit length of threshold. Barni et al.'s [4] performance on the ECG dataset ($d = 4, m = 6$ and $n = 4$) was 6.8 s in computation (without network communication) and 0.1122 MB in communication size. The performance of our method with the same parameters was about 8.85 times better (0.768 s) than [4] in terms of computational cost,

although our method incurred slightly more cost in communication (0.156 MB). Additionally, the security level in our experimental setting was higher than that of [4,6]. (They conducted experiment at a 80 bit security.)

We also conducted experiments on DTs trained using several real datasets used by conventional methods [5,30,35]. Even compared to the methods specialized in DT, the experimental results showed that the performance of our method exceeds that of Bost et al.'s method [5] and is almost equivalent to those of Wu et al.'s and Tai et al.'s methods [30,35]. Additionally, our protocol had an advantage on deep decision trees over [35] whose complexity is exponential to d. When $d = 17, m = 58$, and $n = 57$, our method achieved about a 4 fold faster runtime. The methods specialized in DTs can be used for BDGs by transforming a BDG into an equivalent DT. Their computational cost increases along with the increase of redundant nodes and edges incurred by the transformation. Therefore, while the state-of-the-art method by Tai et al. [30] performed slightly better than our method for DT evaluation, its runtime became worse than our method's runtime when it is tested on complex BDGs.

5 Conclusion

We proposed an efficient protocol for evaluating BDGs, which was designed by AHE and did not use heavy cryptographic primitives, such as fully homomorphic encryption. The protocol obliviously evaluated a look-up vector that was constructed based on GLOUDS to achieve linear time and communication complexities. We also proposed a design principle for the look-up vector to further reduce the vector size. The results obtained from the experiments indicated that the actual runtime and communication size were well concordant with theoretical complexities and that the runtime of our method was an order of magnitude faster than that in the previous approaches [4,6]. We also confirmed that our method was even faster for the DT evaluations compared to a previous approach that specialized in DT [35] when the tree was deep. Our method demonstrated a runtime in an experiment with BDGs that was faster than the state-of-the-art method of DT evaluation [30] that took advantage of the fact that a graph with information equivalent to that in a tree was much more compact than the tree. These results confirmed the efficiency of our method, and we also hope that it will contribute to secure utilization of valuable classifiers that aggregate knowledge extracted from abundant data resources. Another remarkable feature of our protocol is that it directly simulates step-by-step graph traversal, whereas other efficient methods [30,35] reformulate the graph traversal as an evaluation of polynomial equations. By using an additional look-up vector, our protocol enables traversals both to ascendant and to descendant. Such a feature could be useful for various applications that require more complex graph traversal (i.e. searching on DFA).

Acknowledgements. A part of this work is supported by Okawa Foundation Research Grant and JST CREST grant numbers: JPMJCR1503, JPMJCR1688, JPMJCR14D6.

References

1. Amazon machine learning - predictive analytics with aws (2017), https://aws.amazon.com/aml/
2. Google cloud machine learning at scale — google cloud platform (2017), https://cloud.google.com/products/machine-learning/
3. Machine Learning - Microsoft Azure (2017), https://azure.microsoft.com
4. Barni, M., Failla, P., Kolesnikov, V., Lazzeretti, R., Sadeghi, A.R., Schneider, T.: Secure evaluation of private linear branching programs with medical applications. In: Michael Backes, P.N. (ed.) ESORICS. LNCS, vol. 5789, pp. 424–439. Springer, Heidelberg (2009)
5. Bost, R., Popa, R., Tu, S., Goldwasser, S.: Machine Learning Classification over Encrypted Data. In: NDSS. pp. 1–14 (2015)
6. Brickell, J., Porter, D.E., Shmatikov, V., Witchel, E.: Privacy-preserving remote diagnostics. In: CCS. pp. 498–507. ACM Press, New York (2007)
7. Bryant, R.E.: Graph-based algorithms for boolean function manipulation. IEEE TC **100**(8), 677–691 (1986)
8. Carmit Hazay, Y.L.: Efficient Secure Two-Party Protocols. Information Security and Cryptography, Springer, Heidelberg, 1st edn. (2010)
9. Certicom Research: Standards for Efficient Cryptography 2 (SEC 2) : Recommended Elliptic Curve Domain Parameters (2010), http://www.secg.org/sec2-v2.pdf
10. Cock, M.D., Dowsley, R., Horst, C., Katti, R., Nascimento, A.C.A., Newman, S.C., Poon, W.S.: Efficient and private scoring of decision trees, support vector machines and logistic regression models based on pre-computation. Cryptology ePrint Archive, Report 2016/736 (2016), https://eprint.iacr.org/2016/736
11. Cramer, R., Gennaro, R., Schoenmakers, B.: A secure and optimally efficient multi-authority election scheme. In: Fumy, W. (ed.) EUROCRYPT. LNCS, vol. 1233, pp. 103–118. Springer, Heidelberg (1997)
12. Damgård, I., Geisler, M., Krøigaard, M.: Efficient and Secure Comparison for On-Line Auctions. In: Josef Pieprzyk, Hossein Ghodosi, E.D. (ed.) ACISP. LNCS, vol. 4586, pp. 416–430. Springer, Heidelberg (2007)
13. Dowlin, N., Gilad-Bachrach, R., Laine, K., Lauter, K., Naehrig, M., Wernsing, J.: Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In: ICML. pp. 201–210. JMLR (2016)
14. Fischer, J., Peters, D.: GLOUDS: Representing tree-like graphs. Discrete Algorithms **36**, 39 – 49 (2016)
15. Fredrikson, M., Jha, S., Ristenpart, T.: Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In: CCS. pp. 1322–1333. ACM Press, New York (2015)
16. Fredrikson, M., Lantz, E., Jha, S., Lin, S.: Privacy in Pharmacogenetics: An End-to-End Case Study of Personalized Warfarin Dosing. In: USENIX. pp. 17–32 (2014)
17. Huang, C.L., Chen, M.C., Wang, C.J.: Credit scoring with a data mining approach based on support vector machines. Expert Systems with Applications **33**(4), 847–856 (2007)
18. Jacobson, G.: Space-efficient static trees and graphs. In: FOCS. pp. 549–554. IEEE Press, New York (1989)
19. Kohavi, R., Li, C.H.: Oblivious decision trees graphs and top down pruning. In: IJCAI. pp. 1071–1077. Morgan Kaufmann, San Francisco (1995)

20. Lavecchia, A.: Machine-learning approaches in drug discovery: methods and applications. Drug Discovery Today **20**(3), 318–331 (2015)
21. Liu, J., Juuti, M., Lu, Y., Asokan, N.: Oblivious Neural Network Predictions via MiniONN Transformations. In: CCS. pp. 619–631. ACM Press, New York (2017)
22. Madabhushi, A., Lee, G.: Image analysis and machine learning in digital pathology: Challenges and opportunities. Medical Image Analysis **33**, 170 – 175 (2016)
23. Meinel, C.: Modified branching programs and their computational power, LNCS, vol. 370. Springer, Heidelberg (1989)
24. Mitsunari, S.: C++ library implementing elliptic curve elgamal crypto system, `https://github.com/herumi/mcl`
25. Mohassel, P., Niksefat, S.: Oblivious Decision Programs from Oblivious Transfer: Efficient Reductions. In: Keromytis, A.D. (ed.) FC. LNCS, vol. 7397, pp. 269–284. Springer, Heidelberg (2012)
26. Oliveira, A.L., Sangiovanni-Vincentelli, A.: Using the minimum description length principle to infer reduced ordered decision graphs. Machine Learning **25**(1), 23–50 (1996)
27. Oliver, J.J.: Decision graphs: an extension of decision trees. Technical report, Monash University, Department of Computer Science (1992)
28. Shimizu, K., Nuida, K., Ratsch, G.: Efficient privacy-preserving string search and an application in genomics. Bioinformatics **32**(11), 1652–1661 (2016)
29. Shotton, J., Sharp, T., Kohli, P., Nowozin, S., Winn, J., Criminisi, A.: Decision Jungles: Compact and Rich Models for Classification. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) NIPS. pp. 234–242. Curran Associates, Inc., New York (2013)
30. Tai, R.K., Ma, J.P., Zhao, Y., Chow, S.S.: Privacy-preserving decision trees evaluation via linear functions. In: ESORICS. LNCS, vol. 10493, pp. 494–512. Springer, Heidelberg (2017)
31. Troncoso-Pastoriza, J.R., Katzenbeisser, S., Celik, M.: Privacy preserving error resilient dna searching through oblivious automata. In: CCS. pp. 519–528. ACM Press, New York (2007)
32. Vaidya, J., Yu, H., Jiang, X.: Privacy-preserving svm classification. Knowledge and Information Systems **14**(2), 161–178 (2008)
33. Verizon: IP Latency Statistics — Verizon Enterprise Solutions (2017), `http://www.verizonenterprise.com/about/network/latency/`
34. Veugen, T.: Improving the DGK comparison protocol. In: WIFS. pp. 49–54. IEEE Press, New York (2012)
35. Wu, D.J., Feng, T., Naehrig, M., Lauter, K.: Privately Evaluating Decision Trees and Random Forests. PoPETS (4), 1–21 (2016)
36. Yao, A.C.C.: How to generate and exchange secrets. In: FOCS. pp. 162–167. IEEE Press, New York (1986)