

Sylvain Chatel*, Apostolos Pyrgelis, Juan Ramón Troncoso-Pastoriza, and Jean-Pierre Hubaux

SoK: Privacy-Preserving Collaborative Tree-based Model Learning

Abstract: Tree-based models are among the most efficient machine learning techniques for data mining nowadays due to their accuracy, interpretability, and simplicity. The recent orthogonal needs for more data and privacy protection call for collaborative privacy-preserving solutions. In this work, we survey the literature on distributed and privacy-preserving training of tree-based models and we systematize its knowledge based on four axes: the learning algorithm, the collaborative model, the protection mechanism, and the threat model. We use this to identify the strengths and limitations of these works and provide for the first time a framework analyzing the information leakage occurring in distributed tree-based model learning.

Keywords: decision-tree induction, collaborative learning, privacy-preserving protocols, leakage analysis

DOI 10.2478/popets-2021-0043

Received 2020-11-30; revised 2021-03-15; accepted 2021-03-16.

1 Introduction

Tree-based models are currently among the most powerful data-mining methods. They are widely used in the industry [1, 129] and in machine-learning competitions [11, 150]. These algorithms perform very well for tabular problems with numerical and categorical data, which places them in the top ten of machine-learning methods of 2017 [141] with numerous applications: fraud detection [55], medical diagnosis [9], and stock trading [124]. An important feature of tree-based models is *interpretability*, as it makes them an ideal candidate for the interpretable and explainable machine-learning quest of the last few decades [39, 66, 131]. Interpretability is related to the comprehensibility of a model, *i.e.*, to what extent the end-user is able to com-

prehend the model's learning rationale and to verify the soundness of its decisions. This is invaluable in several domains, *e.g.*, medicine and finance, where black-box machine-learning approaches are not acceptable. Indeed, when conducting medical experiments, researchers seek to identify the key factors that affect their outcome, not just the best predictive model on some data. As pointed out by Freitas [66], the interpretability of decision trees makes them reliable, facilitates analysis, and orients future research by identifying points of interest.

Collaborative learning refers to the setting where a group of entities seeks to train a model on their joint data. Collaborative (also known as *federated* [126]) learning has received much traction, due to its applicability in settings where the data is scarce and distributed among multiple entities. For instance, in medical research, a rare disease is possibly not well represented in the patient data of one sole institution; hence the need for data sharing across diverse institutions to create a generalizable model with strong predictive performance.

However, collaborative learning raises privacy and security issues. The training data, *e.g.*, a hospital's patient medical records, is sensitive and cannot, without appropriate safeguards, be shared with other institutions. This is also reflected by the introduction of strict privacy regulations, such as HIPAA [161] and GDPR [50], that forbid data sharing without proper anonymization or pseudonymization procedures. In particular, protecting data used in collaborative machine-learning pipelines is critical, as recent research introduces various successful privacy attacks [81, 89, 128, 135]. Any information exchanged while jointly training a machine-learning model can break the privacy of the training data or the resulting model.

In this work, we perform a cross-field systematization of knowledge on privacy-preserving collaborative training of tree-based models such as decision-trees, random forests, and boosting. Our systematization is based on four axes: the learning algorithm, the collaborative model, the protection mechanism, and the threat model. Our study emphasizes the usage of privacy-enhancing technologies, showing their strengths and limitations. We find that tensions arise as the learning, distributed environment, and privacy protections introduce new

*Corresponding Author: Sylvain Chatel: Laboratory for Data Security – EPFL, E-mail: sylvain.chatel@epfl.ch
Apostolos Pyrgelis, Juan Ramón Troncoso-Pastoriza, Jean-Pierre Hubaux: Laboratory for Data Security – EPFL, E-mail: first.last@epfl.ch

constraints. Elegant and efficient solutions exist but often at the cost of some information leakage, and the few end-to-end protected solutions are not amenable to all scenarios. Therefore, we also provide a framework that identifies the information leakage occurring during the collaborative training of tree-based models. Our systematization enables us to identify limitations such as relaxed threat models and the lack of end-to-end confidentiality, and overall highlights avenues for future work.

The remainder of the document is structured as follows. In §2, we present an overview of our systematization methodology. In §3, we provide background information on tree-based model learning. In §4, we expand on the learning algorithms and, in §5, on the types of collaborative settings. In §6, we present the privacy-preserving mechanisms used in the literature and in §7 the considered threat models. In §8, we present our leakage analysis framework. In §9, we give an overview of the evaluation conducted in the literature. Finally, in §10, we discuss open challenges and we conclude in §11.

1.1 Related Work

There exist a few works similar to ours. While several works surveyed privacy-preserving data mining [120, 170] or decision-trees [164], these works only sketched the collaborative and learning settings and did not delve into the challenges they induce. Recent works focused on federated learning [61] and the security and privacy issues of machine learning [139], but none from the perspective of tree-based models. The literature on decision-tree classification with differential privacy was investigated [61] but only in the centralized settings, where one entity holds the data (see §5). Similarly, a recent work systematized the knowledge on privacy-preserving decision-tree inference [98]. Our work aims at bridging these gaps by systematizing the topic of privacy-preserving collaborative tree-based models focusing on the specific challenges induced by tree-learning in the distributed setting.

1.2 Terminology

Let a *party* be an entity owning a *local* dataset consisting of samples with multiple features. A party seeks to participate in a tree-based model induction process with other parties. A *miner* is an entity that performs computations. It either conducts the model training on the data owned by the parties or simply assists with inter-

mediate computations. An *aggregator* is an entity that combines, during the learning process, information from multiple parties. We also employ these definitions to account for the non-colluding servers model employed in several works [5, 53, 54, 107, 109, 113, 123]. A *collective* is a group of parties interested in training a tree-based machine-learning model on their joint *global* dataset.

2 Scope and Method

We systematize the research efforts on privacy-preserving collaborative learning of decision-tree models in a thorough and structured manner. Overall, our focus is on the perspective of privacy-enhancing technologies (PETs). Hence, our goal is to understand their use for tree-based model induction algorithms, their application to the distributed setting, their trust assumptions, the challenges that they are confronted with, and their limitations and bottlenecks. Thus, we survey the current literature and analyze it from various viewpoints. In this section, we describe the methodology that we employed when searching the literature and the systematization approach that we devised to classify the relevant works.

Search Methodology. We used Google Scholar [3], Microsoft Academic [4], and DBLP [2], to discover works related to privacy-preserving collaborative tree-based model learning: Our search results comprised 73 papers from a wide range of research communities (see Appendix B). We cross-referenced each paper to discover additional relevant works. Overall, our search resulted in 103 papers about privacy-preserving collaborative learning of decision trees that we analyzed.

Systematization Methodology. To classify and organize these works in a structured manner, we devise a systematization methodology that enables us to characterize their approaches on collaborative and privacy-preserving tree-based model induction. Our method takes into account four systematization axes that we briefly describe here:

Learning Algorithm (see §4). This refers to the techniques used for the tree-based model learning. These include the machine-learning task, the data type, the training algorithm, and the underlying quality metric.

Collaborative Model (see §5). This axis is related to the entities involved in the training of the tree-based model, the computation, communication, and data model assumed: the actors involved, their roles, how they interact, and how the data is distributed among them.

Protection Mechanism (see §6). The protection mechanism refers to the privacy-enhancing technologies employed to protect the different components that interact during the tree-based learning, *i.e.*, the training data, the intermediate values, and the final model weights.

Threat Model (see §7). This systematization criterion is related to the definition of the capabilities and objectives of the adversaries that each work considers. It refers to the assumptions imposed to reach a target protection level for each component of the model learning.

3 Decision-Tree Algorithms

We present various tree-based learning algorithms. Hereinafter, we adopt the following notation. We denote by \mathcal{D} the global dataset of n training samples (\mathbf{x}, y) , where \mathbf{x} consists of d features and $y \in C$ denotes its associated label, with C the set of possible classes. For simplicity, we describe the notation and algorithms for classification tasks. Let \mathcal{A} be the attribute set defined as $\mathcal{A} = \{A_1, \dots, A_d\}$. Each component $A_k \in \mathcal{A}$ is a set indicating the possible values obtained by the k -th feature of a sample. With a slight notation abuse, $A_k = v$ indicates that the k -th feature takes the value $v \in A_k$. For any set S , $|S|$ denotes its cardinality.

3.1 Background on Decision Trees

Decision-tree algorithms are tree-based supervised learning techniques used for classification or regression tasks. The structure of a tree can be seen as a directed acyclic graph that consists of nodes and edges. The *root* node has no parent edges, and subsequent nodes are called *internal*. If an internal node has no children, it is called a *leaf*. During learning, let D_i be the part of the dataset that reaches a node i and that is split among its children. This decision is made depending on D_i 's value for a specific feature called a *split point*. The learning (or induction) process of a tree determines the splitting feature and split point for every tree node and results in what is known as the tree's *structure*.

3.2 Greedy Algorithms

Classic decision-tree induction techniques rely on greedy algorithms [24, 144]. These algorithms follow a top-down approach, *i.e.*, they build the tree from the root

to the leaves. At each node, they attempt to find the optimal “short-term” decision. Thus, for each node, the learning problem is reduced to maximizing an objective function that depends on the algorithm.

Quality Metrics. Among the most popular metrics for the objective function are the entropy-based information gain, the gain ratio, and the Gini index.

Entropy-Based Information Gain. The entropy of the i -th node over the class labels is defined as

$$H_C(D_i) = - \sum_{c \in C} \frac{|D_i^c|}{|D_i|} \cdot \log_2 \left(\frac{|D_i^c|}{|D_i|} \right), \quad (1)$$

with $|D_i^c|$ the number of samples in D_i with class label $c \in C$. The best split is defined as the partitioning of the dataset D_i along the k -th feature that maximizes the information gain $\text{Gain}(k, D_i) = H_C(D_i) - H_C(D_i|k)$, where $H_C(D_i|k)$ is the conditional information given by the k -th feature.

Information-Gain Ratio. For the k -th feature, this metric is the ratio between the gain $\text{Gain}(k, D_i)$ and the information value of the split defined as $H_{A_k}(D_i)$. This metric can be extended to a continuous feature a by splitting its space in two parts (*i.e.*, $a \leq v$ and $a > v$) [181].

Gini Index. Informally, it corresponds to the probability of incorrectly classifying a sample when randomly picking the label following the class value distribution for a specific feature. The Gini impurity metric is defined by replacing the entropy $H_C(\cdot)$ by the Gini function defined as $\text{Gini}(D_i) = 1 - \sum_{c \in C} (|D_i^c|/|D_i|)^2$. This metric is more arithmetically-friendly than the previous two, as it can be rewritten with only additions and multiplications. Similarly to the information gain, the best split of D_i is along the feature that maximizes the Gain value.

Other Splitting Techniques. Some works explore alternatives to these classic splitting techniques. The median-based splitting criterion favours a balanced distribution of the data across leaves [34]. The max operator corresponds to the mis-classification rate if the majority class is returned for a specific feature value [68, 114, 140, 192].

Induction Algorithms. There exist several algorithms that use the above metrics to build trees. The ID3 algorithm [144] builds classification trees over categorical data. Following a top-down approach, and recursively at each node, the best split is computed on the dataset that reaches the node and the available features by using the information gain metric. All possible values of the selected feature are added as children, and the dataset is split accordingly. The process is repeated for every child until all dataset samples belong to the

same class or until the feature set is empty (the majority class is returned). The CART algorithm [24] accounts for both regression and classification tasks by using the Gini impurity metric. After generating the tree, it prunes it into a smaller one. The C4.5 algorithm [145] improves ID3 by using the gain ratio to handle both numerical and categorical data.

3.3 Random Algorithms

Contrary to the greedy algorithms above, random approaches generate a tree at random instead of using heuristics: For each node, a random feature is picked from the set of available ones, the tree is split based on all its possible values, and the selected feature is removed from the set. The structure of the tree is determined randomly beforehand, and the training data is used to prune the tree and label the leaves. This method is known as completely random trees (CRTs) [51]. Geurts *et al.* [74] proposed extremely random trees (ExRTs) similar to CRTs but each node evaluates K random splits and selects the one returning the best information gain.

3.4 Random Forests

Breiman proposed the combination of multiple decision trees to form a random forest (RF) [23]. A forest is an ensemble of classification (or regression) trees trained independently. Each tree is trained with one of the aforementioned induction algorithms, using a set of bootstrap training samples and a random subset of features. Forest predictions are made by aggregating all the trees evaluations. We denote by CRF a RF made of CRTs.

3.5 Boosting

Adaptive Boosting (AdaBoost) [67]. AdaBoost works sequentially by adding simple models (classifiers or regressors) to the ensemble model with a certain weight: The final model is a weighted aggregate of these weak learners whose weight is decreasing according to their error. At each iteration, a model is trained focusing on mis-classified samples from the previous iteration.

Gradient-Boosting Decision Tree (GBDT) [69]. GBDT also follows a sequential approach: At each iteration, the updated model learns from the previous one. This is achieved by training a simple tree model

on the residuals, *i.e.*, the difference between the observation and the output, of the previous tree. In GBDT, residuals for a datapoint (\mathbf{x}, y) are approximated at each iteration as the negative gradient of an objective function $L(y, f(\cdot))$ evaluated on \mathbf{x} with $f(\cdot)$ the sum of classifiers from previous iterations. We refer the reader to Friedman’s work [69] for further details.

XGBoost [31]. This is an improvement of GBDT and currently the state-of-the-art tree-based algorithm [150]. Its main characteristic is the use of the second-order derivatives and a better regularization. XGBoost creates an ensemble of K CARTs and its objective function is $L = \sum_{i=1}^n l(\hat{y}_i, y_i) + \sum_m \Omega(f_m)$, where $l(\cdot, \cdot)$ is a convex loss-function measuring the difference between the target label y_i and the prediction \hat{y}_i , f_m the m -th decision tree with $m \in \{1, \dots, K\}$, and $\Omega(f)$ a regularisation term. At each iteration t , the simplified objective is $\tilde{L}^{(t)} = \sum_{i=1}^n [g_i f_t(\mathbf{x}_i) + \frac{1}{2} f_t^2(\mathbf{x}_i)] + \Omega(f_t)$ where f_t is the t -th iteration CART tree, g_i and h_i are the first- and second-order derivative of $l(y_i, \cdot)$, respectively. Then, representing a tree f_k as a fixed structure with T leaves of weights w , λ and γ parameters, and for a partition I of data reaching leaf j , the optimal weight of j is $w_j^* = -(\sum_{i \in I_j} g_i) / (\sum_{i \in I_j} g_i + \lambda)$. The optimal tree structure is obtained with a greedy algorithm that finds the split maximizing the objective L_{split} . Denoting by I_L and I_R the instances of the left and right nodes, respectively, (*i.e.*, $I = I_L \cup I_R$), the split objective is

$$L_{\text{split}}(I) = \frac{1}{2} \left(\frac{(\sum_{i \in I_L} g_i)^2}{\sum_{i \in I_L} h_i + \lambda} + \frac{(\sum_{i \in I_R} g_i)^2}{\sum_{i \in I_R} h_i + \lambda} - \frac{(\sum_{i \in I} g_i)^2}{\sum_{i \in I} h_i + \lambda} \right) - \gamma. \quad (2)$$

4 Learning Algorithm

As presented in §3, there exist numerous algorithms that can be used to learn tree-based models. Among these, ID3 is predominantly considered for the case of privacy-preserving collaborative decision trees, with *classification* being the most popular task (see Table 1). Hence, most of the surveyed works only consider *categorical* data, whereas some of them propose the discretization of continuous data to account for *numerical* attributes [62]. We classify as “ensemble” any work that creates a forest rather than single trees. For optimization reasons, several works modify the quality metric used by the learning algorithm [5, 17, 34, 38, 53, 54, 112, 114, 122, 149, 178, 183] (see §6). Overall in the literature, we observe a wide range of combinations on tree-types, tasks, data, and algorithms, as illustrated

Table 1. Learning algorithms in the surveyed literature. (Ens.: Ensemble, Clas.:Classification, Reg.:Regression, Num.:Numerical, Cat.:Categorical).

Reference	Task				Data	Algorithm
	Ens.	Clas.	Reg.	Num.	Cat.	
[179, 180]	•	•	•	•	•	CART
[37]	•	•	•	•	•	CART
[34]	•	•	•	•	•	CART-like
[5]	•	•	•	•	•	C4.5-like
[168]	•	•	•	•	•	CRT
[116]	•	•	•	•	•	ExRT
[119]	•	•	•	•	•	ExRT
[62]	•	•	•	•	•	CART
[15, 17, 71, 85, 96, 110, 152, 154, 158, 181]	•	•	•	•	•	C4.5
[41]	•	•	•	•	•	CART
[172, 192]	•	•	•	•	•	CART-like
[156]	•	•	•	•	•	CRT
[68, 171]	•	•	•	•	•	ID3/C4.5
[112]	•	•	•	•	•	C45
[7, 8, 27]	•	•	•	•	•	CART
[9]	•	•	•	•	•	C4.5
[16, 26, 79, 183]	•	•	•	•	•	CART
[95, 108, 177]	•	•	•	•	•	CRT
[84]	•	•	•	•	•	ID3/CRT
[20]	•	•	•	•	•	CRF
[190]	•	•	•	•	•	GBDT
[75]	•	•	•	•	•	RF
[117]	•	•	•	•	•	RF/CART
[104, 105]	•	•	•	•	•	GBDT
[32, 101, 102, 118, 162, 175, 184]	•	•	•	•	•	XGBoost
[12, 60]	•	•	•	•	•	CRF
[83, 103, 140, 146]	•	•	•	•	•	RF
[70, 173]	•	•	•	•	•	AdaBoost
[123]	•	•	•	•	•	RF
[55, 56]	•	•	•	•	•	XGBoost
[137]	•	•	•	•	•	Boosted
[59]	•	•	•	•	•	CRF
[188]	•	•	•	•	•	ExRT
[10]	•	•	•	•	•	RF
[19, 38, 42, 48, 52–54, 63, 72, 73, 76, 80, 90, 100, 107, 109, 111, 113, 114, 121, 122, 147, 149, 153, 157, 160, 163, 166, 167, 177, 182, 186, 187]	•	•	•	•	•	ID3

on Table 1. We remark that the learning algorithm imposes constraints on designing a privacy-preserving and collaborative solution as we will detail this issue in subsequent sections.

5 Collaborative Model

We propose a systematization of the literature on privacy-preserving collaborative tree-based model learning based on their collaborative model. In distributed settings, it is crucial to understand which entities are involved, their role, and how they interact. We distinguish two major aspects of the collaborative model: (i) the computation and communication model and (ii) the

data distribution model. The former considers the entities and their interactions, whereas the latter describes how the global dataset is partitioned.

5.1 Computation and Communication Model

We consider three types of entities that participate in the distributed learning: parties, miner(s), and aggregator (see §1.2). Although some works rely on additional (external) entities (*e.g.*, a trusted third-party, a public-key infrastructure), here we omit them as they do not directly participate in the learning process. We present the different collaborative models considered in the literature on distributed privacy-preserving tree-based model induction. The different categorization corresponds to where the training data is available and where the bulk of the computation is executed.

Central and Offloading Model. In this case, a miner has access to all the parties’ data and performs the training. This model covers works that are not initially envisioned for the distributed setting, *e.g.*, [19, 26, 34, 62, 68, 84, 105, 114, 140, 183, 192], as well as works under the *offloading* category where the parties explicitly communicate their dataset to the miner, *e.g.*, [7, 8, 12, 26, 49, 63, 123]. Brickell and Shmatikov propose a model where a miner creates a private model on a single remote secret database [26]. For privacy, Abspoel *et al.* assume parties offload their data by sharing it across three non-colluding servers [5]. The offloading to only two servers is also possible [123].

Aggregator Model. Contrary to the previous model, in this one, any party that takes part in the learning process communicates with an aggregator. The latter coordinates the training by obtaining intermediate values computed locally by the parties on their data. The aggregator combines these values and publishes the result such that the learning process can continue. The aggregator operates differently, depending on the learning task. For example, it issues count queries to the parties to compute the information gain for the ID3 algorithm (*e.g.*, [163]) or selects a subset of features to be considered by parties (*e.g.*, [116]). In another approach, trees are learned locally and sent encrypted to the aggregator (*e.g.*, [75]). The aggregator uses all the locally trained encrypted trees to compute the encrypted prediction. But, experiments show a performance drop in the model accuracy [75]. A similar approach is also used for distributed AdaBoost [106, 180].

Fully Distributed (FD) Model. This scenario involves only members of the collective that hold local data. In this case, each party communicates with all the others. Although initially considered for the setting of two parties by Lindell and Pinkas [111] and subsequent works [41, 55, 56, 76, 160, 181], it was later extended to arbitrary number of parties [38, 48, 70, 100, 102, 104, 106, 149, 168, 179, 182]. In this setting, some parties might be assigned specific tasks. We refer to systems with consistent parties’ roles during learning as *leader-based* [32, 80, 153, 156, 186, 187], whereas those with temporary ones as *sequential* [190]: *e.g.*, the local learning of a model before communicating it to the next member of the collective.

5.2 Data-Distribution Model

The global dataset can be partitioned in different ways among the entities that participate in the tree-based model-learning process:

No Partitioning. The whole data is centralised in one unique site and the learning is executed on it. It is directly related to the central and offloading models presented in §5.1 and often a comparison baseline. Although it requires parties to communicate their data to a remote server, no subsequent rounds of communication are needed as the learning operates as if the data was never distributed.

Horizontal Partitioning. All the parties possess their own data samples that share the same feature space. The collaborative learning incurs communication overhead as the parties need to exchange intermediate results computed on their local data. For instance, for the ID3 algorithm, the collective needs to compute the information gain for every possible split: This requires sharing the counts of the data points at each split. This becomes communication heavy, with an increasing number of features and classes, hence multiple works modify the learning algorithm to a randomized one that is more collaboration-friendly; indeed, only the final leaf counts need to be computed on the distributed data.

Vertical Partitioning. Conversely to the previous case, with vertical partitioning, the parties share a different feature space for the same individual samples of the global dataset. This can, for instance, represent a distributed database containing the data of a common set of customers across multiple financial institutions. The class label is known by all the parties or by only one. The challenge in this setting is to find the feature on

which to split the tree. With greedy algorithms, to find the feature with the best gain, each party can locally compute the gain, and then communicate with others.

Arbitrary Partitioning. In this case, both horizontal and vertical partitioning are present. It is often referred to as an arbitrary (or mixed) partitioning of the data. Only a few works present a distributed solution for this case by using special data representation [80] or by adding extra rounds of communication [38, 100, 122, 168].

5.3 Summary

Table 2 presents the collaborative models of the surveyed works. The vast majority of works consider the aggregator or offloading model. Works in the central model can trivially handle both data partitioning tasks, as they eventually gather the data in one place. The fully distributed case is by far the most challenging one. However, most of those solutions focus on the two-party case. Vertical and horizontal partitioning have, respectively, their own limitations: The former relies on local best attributes but often needs public class labels, and the latter requires distributed computations, introducing computation and communication overhead. Few works acknowledge the challenges introduced by the local data following different statistical distributions [61, 103, 104, 106, 190]. Although this is a non-issue for greedy algorithms in which the distributed algorithm works as if the data was centralised, it can become cumbersome for randomized and boosted algorithms [103]. Moreover, the notion of availability or dropout of the parties is often overlooked with only one work that provides a solution for parties to leave the collective and that removes their impact on the training [119].

6 Protection Mechanism

We review the types of privacy-enhancing technologies (PETs) employed to ensure confidentiality of data during the collaborative tree induction process. We identify five categories of PETs: (a) input randomization (§6.1), (b) differential privacy-based solutions (§6.2), (c) cryptographic approaches (§6.3), (d) hardware-based solutions (§6.4), and (e) hybrid solutions that combine the above (§6.5).

Table 2. Different collaborative models in the literature.

Reference	Coll. Model	Members		Data Model			
		Parties	Aggregator	Miner	Vertical	Horizontal	Both
[38, 100, 122, 168]	FD	●					●
[71, 157, 167, 179]	FD	●			●		
[17, 48, 70, 102, 149, 182]	FD	●				●	
[95]	FD (2 parties)	●					●
[41, 55, 56, 76, 96, 160]	FD (2 parties)	●			●		
[111, 121, 181]	FD (2 parties)	●				●	
[80]	Leader	●					●
[32, 37, 153, 166, 186]	Leader	●			●		
[72, 104, 156, 187]	Leader	●				●	
[190]	Sequential	●				●	
[110, 162]	Aggregator	●	●				●
[52, 117, 119, 152, 154, 172]	Aggregator	●	●		●		
[9, 10, 73, 75, 101, 103, 108, 116, 118, 163, 175, 180, 184]	Aggregator	●				●	
[5, 53, 113, 123]	Offloading	●	●	●			●
[54, 107, 109]	Offloading	●	●	●		●	
[5, 7, 8, 12, 15, 27, 42, 63, 85, 90, 112, 130, 147, 158, 173]	Offloading	●		●			●
[16, 19, 20, 26, 34, 59, 60, 62, 68, 79, 83, 84, 105, 114, 137, 140, 146, 171, 183, 188, 192]	Central	●		●			●

6.1 Input Perturbation and Randomization

These techniques generate a surrogate dataset by perturbing the original one to protect its confidentiality; this dataset is used by the miner to train the tree-based model. Thus, it is predominantly employed in the offloading collaborative model, *e.g.*, [7, 27, 40, 49, 63, 85, 112]. Typically, perturbation techniques discretize or add noise to each sensitive attribute of the dataset [7], or swap the dataset with a surrogate one that has the same probability distribution [49]. However, Kargupta *et al.* [93] showed that noise addition does not prevent the reconstruction of patterns in the data: They propose a filtering technique that reconstructs the original dataset. Whereas other random substitution techniques are immune to such attacks [40], they hamper utility as the mining is performed on an approximation of the original data [63]. Some works reconstruct a surrogate dataset from the sanitized original database [90], whereas others assume that the miner obtains the tree computed on the original data before adding noise [85]. Finally, some studies modify the learning algorithm to handle the perturbed datasets [112].

Although this line of work does not necessarily target the distributed scenario, it can be trivially extended to it. Overall, input perturbation and randomization techniques are ad-hoc methods that obfuscate the train-

ing dataset with limited privacy guarantees. So, more recent works rely on differential privacy to address this lack of formalization.

6.2 Differential Privacy Based Solutions

Introduced by Cynthia Dwork [44], differential privacy (DP) is a privacy definition for sharing information with respect to queries on a database. Informally, it guarantees that the change of a single data record does not modify the query’s output significantly. For a query f , this is achieved by adding noise to its answer; the noise amount is determined by the query’s sensitivity $\Delta(f)$ and the privacy budget ϵ . When designing a DP-based solution for collaborative decision-tree learning, four important aspects should be considered: which entity performs the noise addition, at which training stage, the magnitude of noise required, and the total privacy budget spent for training. Tackling these enables to train and publish trees with DP guarantees.

Central Model. Recall that, in this scenario, a single entity has access to the entire training dataset. The main idea is to inject noise during key learning parts, *e.g.*, for selecting the best feature [68], counting class counts at the leaves [34, 84], or computing gain queries for each feature [62]. More recent works aim to find tighter sensitivity bounds for the training queries or new ways to embed DP [16]. Other approaches *relax* the learning algorithm by replacing information gain with more DP-friendly metrics, *e.g.*, Gini [62] or Max [114]. For instance, the max operator has lower sensitivity than the Gini or information gains thus leading up to higher accuracy on similar datasets and privacy levels [61, 68]. In other works, the learning is adapted using RFs [83, 140, 146] or CRTs [20, 59, 60, 188]. Some works abusively consider that each tree in the forest is independent, to reduce the privacy budget consumption [146]. However, this is circumvented by training each tree on an independent subset of the training data and applying the parallel composition theorem [59, 183]. Overall, works based on the central model consider reasonable privacy budgets (*i.e.*, $\epsilon \in [0.1; 1.0]$) and some even experiment with very low budgets (*e.g.*, $\epsilon = 0.01$ for [16, 20, 59, 114, 137]). We note that the privacy budget configuration directly affects the model’s accuracy; Fletcher *et al.* report accuracy drops of more than 20% when the budget is reduced from 2.0 to 0.2 to obtain stronger privacy guarantees [60]. We refer the interested reader to the survey by Fletcher and Islam on differen-

tially private trees [61], as it exclusively focuses on the central model where the miner has access to the dataset.

Aggregator Model. In this setting, parties apply perturbations to their local intermediate results before sharing them with the potentially untrusted aggregator. Xiang *et al.* [180] present two collaborative algorithms: CART-based random forest and AdaBoost. In their solution, each party builds a local ensemble model by injecting noise into each tree, and the aggregator merges the perturbed models with a weighted sum that depends on each party’s data samples and reported accuracy. Similarly, Liu *et al.* build a differentially private ExRT where the split is randomly selected to reduce the privacy budget consumption [116]. In the work by Li *et al.* [103], each party uses a differentially-private Generative Adversarial Network (GAN) to generate a surrogate “synthetic” dataset. This can be shared with the aggregator; it merges the received datasets and redistributes them to the parties that use them as validation data to select the best local model. Finally, the best local models are shared with the aggregator that combines them into the final global model.

Fully Distributed Model. Zhao *et al.* [190] propose a collaborative system for GBDT. They employ an iterative sequential method where each party locally trains on its data and transfers the resulting tree to the next party. To ensure privacy of the shared model, each party generates DP trees: The split values are sampled using the exponential mechanism [127].

6.3 Cryptographic Solutions

Numerous works employ cryptographic techniques to protect the confidentiality of the tree-based model induction process. The most common cryptographic tools are Secure Multiparty Computation (SMC) (*i.e.*, secret sharing, garbled circuits, and oblivious transfers), homomorphic encryption (HE), encryption as obfuscation, and locality-sensitive hashing. We provide a brief background on these before analyzing their use by the relevant works. Table 4 displays the different cryptographic solutions employed in the literature (see PET column).

Background. Secret sharing (SS) methods distribute a secret into *shares* such that the secret can only be reconstructed by re-combining all the shares (or a subset of them piloted by a threshold). Oblivious transfer (OT) [133], oblivious polynomial evaluation [134], and garbled circuits (GC) [185] are SMC building blocks that enable private secrets exchange,

function computation and, overall, circuit evaluation on private inputs. Generic frameworks (*e.g.*, SPDZ [36, 94] or VIFF [159]) use these techniques and provide an abstraction for SMC supporting arithmetic operations, comparisons, and more. Other useful protocols include private set intersection (PSI) [99, 142, 170], cardinality of intersections [6, 65, 166], and secure scalar product (SSP) [41]; these protocols are at the core of some distributed tree induction algorithms: *i.e.*, counting the number of samples reaching a node using dot product between binary vectors [41] or creating consensus among two parties [70]. Homomorphic encryption (HE) enables computations on ciphertexts without requiring decryption. Depending on the scheme, operations can be linear (LHE) [46, 138] or, with fully homomorphic (FHE) schemes, polynomial ones [22, 33]. HE can reduce the communication overhead of SMC: Instead of being shared, the secret is encrypted and computations are done directly on the ciphertexts. Also, as HE schemes support only limited operations, combining them with SMC in the offloading model enables new functionalities such as divisions or comparisons [115]. Recent schemes merge directly HE and SMC for efficiency [29, 97, 132]. Additionally, HE and SS can be used to keep the final model secret and support oblivious predictions [5, 12, 75, 119, 123, 168]. We now present how SMC and HE techniques are used to protect distributed tree-based model induction by recognizing the constraints imposed by the learning and collaborative model chosen (*e.g.*, required information, communication topology, or data partitioning).

Secure Multiparty Computation. Du and Zhan [41] pioneered the use of SMC for privacy-preserving ID3 training over a vertically distributed dataset among two parties. In their system, an external semi-honest entity generates blinding shares that are used, during the training process, for secure scalar-product operations: The intermediate counts are computed by dot products between binary vectors that represent the constraints needed to reach a specific node. Similarly, Lindell and Pinkas proposed the first algorithm for ID3 induction with horizontally distributed data across two parties [111]. Their algorithm uses GC techniques to obtain the attribute with the minimal conditional entropy and to compute an approximation of the $x \log x$ function required to calculate the information gain. However, it does not scale well with increasing number of parties. To this end, Emekçi *et al.* use Shamir’s secret sharing [151] and propose a new secure-sum protocol to aggregate the counts required for the learning process [48]. Compared

to their previous work [47], they relax the need for a trusted aggregator and include a correctness check to thwart malicious parties that tamper with intermediate results during the computations. In particular, they increase the degree of the random polynomial used for the sharing, and they introduce redundancy in their system, which makes it computationally infeasible for a party to cheat. Nevertheless, their solution enables secure computation of only the gains; the best feature and the data split is performed in cleartext. Ma and Deng [122] reduce the communication and computation overhead of arbitrarily distributed ID3 by replacing the information gain with the Gini impurity metric. Privacy is ensured via secret sharing; multiplication, comparison, and addition operations are achieved with custom protocols executed among the parties. However, a trusted server is required to generate shares of secret blinding values. Gambis *et al.* [70] use secure-sum [91] and PSI [99] protocols to enable distributed AdaBoost. De Hoogh *et al.* [38] opt for a generic framework [159] for SMC by using Shamir secret sharing: They train trees using the ID3 algorithm with the Gini index metric, but their solution supports only categorical data. This limitation was recently addressed by Abspoel *et al.* [5] by using SPDZ [94].

Contrary to perturbation or DP-based approaches, SMC solutions enable almost *exact* learning (to the approximation of non-polynomial operations, *e.g.*, ID3), hence they do not compromise the accuracy of the resulting model. Additionally, a handful of SMC solutions protect the resulting model by keeping it private (*i.e.*, secret-shared) [5, 38]. SMC solutions also accommodate multiple parties. Their main drawback, however, is the introduced computation and communication overhead. Furthermore, to enable specific computations, *e.g.*, multiplications, there is sometimes the need for a trusted setup or a trusted server that generates intermediate values (*e.g.*, blinding shares or Beaver triples) [41, 55, 113, 179].

Homomorphic Encryption. Considered initially for offloading and outsourcing scenarios [8, 12], HE can also be used in collaborative settings by using the appropriate keying material. Most works that employ HE-based approaches use the additive scheme introduced by Paillier [138], its threshold variant [35], or similar multi-party schemes [25]. HE alleviates some limitations of pure SMC approaches, *e.g.*, the communication overhead and the need for a trusted setup. Indeed, with HE, secrets do not need to be shared among all parties, and computations can be executed by a single party

without compromising privacy. On the negative side, HE schemes are limited by the operations allowed on ciphertexts; these schemes might not suffice to execute some tree-based learning algorithms.

Some works employ HE solely for the computation of the gain; the selection of the best attribute is computed in cleartext. Vaidya *et al.* [167] introduce a secure dot product protocol using the Paillier cryptosystem to overcome the two-party limitation of prior work [41]. This protocol improves previous set-intersection algorithms [6, 65, 166], by limiting leakage of unused information. Similar HE-based custom techniques for set-intersection cardinality or scalar product are also employed in the literature [37, 80, 119, 160]. HE is also used for aggregating the local encrypted counts or statistics required to compute the gains [9, 32, 53, 54]. Once aggregated, the result is decrypted and subsequent operations are conducted in cleartext. Several HE-based works opt for the Gini metric as it has an arithmetic representation simpler than the other metrics [53, 54, 149].

Other HE-based works modify the collaborative model. This enables the computation of functions not supported by LHE. In the leader collaborative model, Zhan *et al.* [186, 187], propose solutions with three leading parties: One performs Paillier encryption/decryption and computations, a second generates randomness, and the third is in charge of blindings. This design enables the computation of non-linear functions and comparisons. However, the special-role parties need to be available and to follow the protocol. In the offloading scenario, some works introduce an additional entity [107, 110, 123]. This helps with the computation of non-linear functions and comparisons through multi-party HE protocols among two non-colluding parties [25]. Sometimes parties directly assist with the computations [8]: The gain is computed on cleartext data, and the comparison is replaced by an approximation of the step function.

Finally, other HE-based works relax the learning algorithm and employ CRTs [95, 108, 156, 168]. The training data is used only for updating the leaves' statistics, thus HE can be used to gather these counts [168]. Alternatively, each party creates a CRT or a local tree that is added in the global forest: HE-based consensus can be used to select which trees to retain [108] or the trees are directly shared in encrypted form [75].

Overall, HE is a powerful tool that can reduce the communication overhead of pure SMC solutions. Although it is affected by similar constraints to SMC (*e.g.*, requiring simple arithmetic circuit representations of the computations), HE solutions also introduce new

challenges; the main one is the computation of the best information gain. Finally, we note that the adversarial model influences the choice of the HE scheme and its efficiency: *e.g.*, the widely used Paillier [138] or El-Gamal [46] LHE schemes are not secure against quantum adversaries; recent lattice FHE schemes such as BGV [22] or CKKS [33] alleviate this limitation at the cost of larger ciphertexts. However, we found only a few works that use recent FHE schemes and focus on the offloading model [8, 12]. They need a specific data representation of categorical data (one-hot-encoding or discretization) for equality tests and comparisons [12].

Combining HE and SMC. SMC and HE approaches are complementary. HE can reduce the communication overhead of SMC, whereas the latter supports arithmetic operations (*e.g.*, comparisons) that are inefficient with HE. Thus, many works on privacy-preserving collaborative decision-tree learning combine them to exploit the best of both worlds. For greedy algorithms, HE can be used to compute aggregate intermediate values by combining encrypted local values of the $x \log x$ function, the Gini index, and the gain ratio. The subsequent comparisons to find the best attribute can be done privately using garbled circuits [26, 181]. Kikuchi *et al.* combine SMC [41] and HE [167] approaches, and they design a secure scalar product protocol that incurs low communication costs and does not require a trusted setup [96]. For XGBoost, several works use HE and secret sharing to protect the local intermediate residuals [118, 119, 175]. Liu *et al.* [118] propose an aggregation scheme that, with Shamir secret sharing and Paillier HE, ensures that the aggregator cannot access individual party updates. Each party locally computes gradients and the aggregator derives the score of each party’s split to select in clear the best one. Similarly, Fang *et al.* [55] propose a solution combining additive secret-sharing and HE. Contrary to similar works relying on pure HE [32] or combined with hypervisor-enforced domain isolation and OT [56], their solution maintains every value encrypted or secret-shared, thus does not leak intermediate information. Wu *et al.* [179] port XGBoost to the fully encrypted setting for more than two parties by using a threshold version of Paillier HE and the SPDZ framework [36]. Similarly to the SMC approaches, this work requires the generation of secret shares of random multiplications (Beaver triplets) by a trusted third party. Recently, Liu *et al.* [113] proposed a new offloading solution that could be extended to collaborative scenarios. Using additive HE, data owners offload their data to a cloud. Using additive secret-

sharing and with the help of a computing server, the cloud builds the tree from the encrypted data via tailored secure-counting and comparison protocols. They improve similar works [109, 110], by relaxing the need for parties to be online during the learning.

Encryption as Obfuscation. Similar to perturbation techniques, obfuscation of the sensitive data can be achieved with encryption techniques (*e.g.*, AES) [147, 152, 158]. The encryption can be lifted by the miner once data has been merged [158]. Alternatively, parties offload deterministic encryption of their data, and the learning is done on the ciphertexts as new labels and feature values [147, 152]. But this encryption technique is prone to frequency-analysis attacks [18, 136].

Locality-Sensitive Hashing (LSH). SimFL [104] includes a pre-processing phase where LSH is applied and similar information across parties is grouped without revealing raw data. LSH ensures that similar (resp. dissimilar) instances have equal (resp. different) hash digests with high probability. During training, the gradients of all similar instances are included in the boosting. SimFL improves upon previous work [104, 190] either in terms of accuracy or efficiency. While SimFL is fully distributed, it considers a relaxed threat model where a dishonest party might learn some information about the other parties through inference attacks, but not through their raw data.

6.4 Hardware-Based Solutions

Trusted hardware, *e.g.*, secure enclaves [92, 125], is an alternative solution for private distributed model induction. In particular, a few recent works consider that each party installs a secure enclave at its premises; this enclave is responsible for storing and computing on the sensitive data [101, 102]. Hardware-based solutions impose different trust assumptions and are orthogonal to the aforementioned software-based solutions. Also, recent research shows that secure enclaves are susceptible to side-channel attacks [169, 174].

6.5 Hybrid Solutions

Hybrid solutions combine the various PETs described earlier. For example, Teng *et al.* [160] combine randomization techniques with SMC: Each party’s local dataset is enhanced with perturbed data from others to find the best set of features during tree induction. This leads to a model accuracy better than randomization

approaches at reduced computation costs, compared to SMC. Truex *et al.* [163] propose a hybrid approach for federated learning; it employs DP mechanisms and threshold HE. They apply it to decision-tree learning by using the ID3 algorithm and an aggregator that initiates the root node. They employ threshold additive HE on noisy inputs such that the aggregator decrypts feature counts and class counts values with DP guarantees. Their solution yields a more accurate model than others that employ local DP, as the amount of noise is divided by the number of parties required to “unlock” the HE threshold. Subsequently, Liu *et al.* [118] improved this approach by ensuring that the aggregation is performed correctly. Moreover, combining local DP and secure aggregation with threshold HE [21] is envisioned in both horizontal and vertical data partitioning settings for XGBoost [162]. Finally, Wu *et al.* [179] inject noisy values during training to achieve the guarantees of DP for the output of their SMC-based aggregation system, *i.e.*, the resulting model.

7 Threat Model

We now systematize existing works on privacy-preserving collaborative decision-tree learning based on their threat model (see Table 3). We identify two main threat model categories: (a) honest and semi-honest, and (b) malicious.

Honest and Semi-Honest. An entity is deemed honest if it abides by the protocol and does not try to infer anything from the data exchanged and stored during the protocol. A system operates under an honest model if all the parties involved are honest. While none of the surveyed works considers a fully honest threat model, many solutions rely on at least one honest entity in the system. For example, miners in the central collaborative model are de facto trusted to train the model and to inject the necessary noise to ensure differential privacy guarantees of the output, *e.g.*, [59, 69].

The semi-honest model (also referred to as passive or honest-but-curious) considers that the participants follow the protocol but might try to infer as much information as possible about the other entities’ private data, from the communicated values. This is a typical threat model for solutions based on perturbation and randomization techniques that implicitly consider that parties exchange their data with a semi-honest miner. Indeed, the miner is not trusted with the original datasets but obtains full access to each party’s surrogate

data, *e.g.*, [7, 63]. For works that consider the application of DP in the aggregator collaborative model, the aggregator is often considered semi-honest, hence parties add noise to their local computations before sharing them with the aggregator [116, 180].

Most works using cryptographic techniques consider the semi-honest model for the different entities [32, 41, 55, 70, 100, 111, 113, 166, 181, 182]. Only a handful of these works also consider passive collusions among the different members. For instance, by using SMC frameworks or LHE, several works are secure against a collusion between half of the N involved parties [38, 167]. With threshold encryption or secret sharing other works tolerate up to $N-2$ [70, 80], or even $N-1$ [48, 168], colluding participants. In both works of Truex *et al.* [163] and Wang *et al.* [175], the maximum number of colluding members admissible without damaging privacy is piloted by the threshold defined for the secret sharing scheme. In other works, however, collusion causes a direct loss of privacy: In the work of Du *et al.* [41], the collusion between one party and the third-party assisting with the computation can reveal the other party’s secret data. The revocation mechanism in the work by Liu *et al.* is secure, as long as the revoked member does not collude with the aggregator [119]. Li *et al.* [104] consider that a dishonest party might learn some information about the data of other parties, *e.g.*, local gradients, but not raw data. Theirs is one of the few works that takes into account potential leakage induced by sharing intermediate values during the training process.

Remark. Note that in several works, the privacy is guaranteed by the presence of implicitly trusted third-parties involved in the generation of cryptographic keys [107, 110] and random shares used for the computations (*e.g.*, Beaver triplets for multiplication operations in SMC) [179]. Similarly, hardware-based solutions presume the chip manufacturer is trusted and the attacker does not have access or control over the enclave [101, 102].

Malicious. Also known as active, malicious participants can *actively* cheat and tamper with the protocol by crafting messages with fabricated inputs and by aiming to gain more information about the other entities’ data or to simply disrupt the protocol. A limited amount of works consider resistance against malicious entities. As a party can always tamper with its local training data to perturb the learning, the envisioned malicious model concerns active adversaries who aim to cheat the learning process by performing wrong computations. Emekçi *et al.* [48] rely on Shamir secret-sharing,

along with a technique that verifies the correctness of an aggregate result; by including redundancy to the secret-sharing polynomials, they are able to over-determine the equation system that reveals the result upon solution. They show that with appropriate tuning, it is computationally hard for an adversary to forge a result undetected. Akavia *et al.* [8] provide privacy guarantees against a malicious aggregator that tries to learn as much as possible from the clients’ inputs following any arbitrary attack strategy. Furthermore, a malicious aggregator might be tempted to skip the correct aggregation of intermediate values to obtain some local information. This behaviour is tackled by Liu *et al.* [118] with τ -threshold secret-sharing which ensures that the aggregator only learns the aggregate result over the data of at least τ parties. In InPrivate [190], the malicious parties seek to tamper with the steepest-descent returned. Hence, the design employs a local quality control of other members’ trees: Each party evaluates on its local data the performance of the tree received from the previous member and decides whether to discard it or not. In the work by Wu *et al.* [179], malicious members that deviate from the protocol are considered as an extension of their design: They use zero-knowledge proofs and commitments to prove statements about secret data without disclosing it. In particular, each member of the collective proves that it executed the specified protocol correctly. Similarly, though not considering the malicious model, several works claim that this model can be supported using general techniques such as those presented by Goldreich [77, 78] at the cost of efficiency (in terms of computations and communications) [111, 168]. Recently, Abspoel *et al.* [5] use SPDZ [94] to provide a MPC-based C4.5 algorithm. Active security is achieved assuming honest majority among three non-colluding servers and sacrificing performance.

8 Leakage Analysis

Our analysis of the literature on privacy-preserving collaborative tree-model induction (§6) shows that very few works protect the training process end-to-end. We find that most works ensure the confidentiality of the raw training data, but do not consider the leakage that might occur from computations required for collaborative tree induction, *e.g.*, comparison operations are performed on cleartext data to abide with HE limitations. To this end, we design a framework that analyzes which information is leaked during privacy-preserving collab-

Table 3. Threat models considered in the literature. ○: Honest, ⊙: Honest-but-curious, ●: Malicious, Coll. Res.: Collusion resistant, N number of parties, τ : Secret sharing threshold, u : number of parties knowing the class labels, TH: Trusted Hardware

Reference	Parties	Miner	Aggregator	Coll. Res.
[107, 109, 113, 123]	○	○	○	
[7, 12, 15, 26, 27, 40, 49, 53, 54, 63, 85, 90, 112, 130, 147, 158, 165]	○	○		
[8]	○	●		
[162, 163, 175]	○		○	$N-\tau$
[9, 10, 52, 75, 103, 108, 110, 116, 119, 152, 154, 172, 180, 184]	○		○	
[118]	○		●	
[38]	○			$[N/2]$
[48, 179, 190]	●			$N-1$
[70, 80, 153]	○			$N-2$
[32]	○			$N-u$
[167, 168]	○			$N-\tau$
[17, 37, 41, 42, 55, 56, 71, 72, 95, 96, 100, 104, 111, 121, 122, 149, 156, 157, 160, 166, 181, 182, 186, 187]	○			
[5]	●			$[N/2]$
[101, 102]	TH		TH	
[16, 19, 20, 34, 59, 60, 62, 68, 79, 83, 84, 105, 114, 137, 140, 146, 171, 177, 183, 188, 192]	○	○		

orative tree-based model induction, enabling us to systematize the literature on that aspect. Although several works acknowledge this leakage and even provide an analysis of their solution [32, 55, 70, 75, 118, 167, 168, 179], they do not do so in a systematic way. Our framework is a first step towards a generic systematization of works based on the information leakage. We explain the need for leakage analysis and then present our taxonomy.

8.1 Importance of Leakage Minimization

Overall, revealing computation results during tree-based model induction, *e.g.*, statistics or model updates, produces a potential leakage about the training data. For instance, global statistics refer to values aggregated from several parties involved in the learning. This can include feature counts, class counts, the global feature list, or the party owning the best split. Wu *et al.* [179] describe an attack among colluding parties in the vertical data-partitioning setting that we extend here. Consider that only one party holds the class attribute. Then, colluding parties responsible for successive splits from the root can, with access to the label of the leaf, infer the class of a subset of the training data with high probability. They also propose a second attack that relies on the same principle to yield the feature value. Generally, during tree-based induction in the horizon-

tal data-partitioning setting, aggregate statistics leak information about the global training set (*e.g.*, number of samples, distribution per feature). Similarly, in the vertical case, the chosen metric value for the best local feature is revealed. Even though no concrete attacks that exploit these values have been proposed, this leakage might be non-negligible and should be taken into account by designers of privacy-preserving solutions. Indeed, numerous works have shown that aggregate statistics leak information about individual samples in several settings such as genomics [82], smart metering [28], and location data [143].

Similarly, weight updates might leak information and lead to membership or property-inference attacks [128, 135, 191]. A curious participant could infer the presence of specific samples in other parties' data or properties about the data. This is particularly relevant for works using gradient boosting. Access to local model updates makes the training process vulnerable to updates-leak attacks [148] that aim at inferring information about the training sets from the model updates. Finally, as shown by Hitaj *et al.* [81], a Generative Adversarial Network can be trained on the global aggregate updates to generate prototypical examples of the targeted victim's local dataset. Even though these attacks focus on different learning algorithms, *e.g.*, deep convolutional neural networks, these attacks show that the information communicated during gradient descent contains private data. Fang *et al.* acknowledge and tackle this leakage by using HE and secret sharing to protect the gradients end-to-end [55].

8.2 Leakage Taxonomy

We now present our leakage classification of the literature on privacy-preserving collaborative tree-based models. We use the term *beneficiary* for the entity in the collective who has access to the leak, *i.e.*, who is able to infer some information. This can be the miner, the aggregator, the parties themselves or a combination of these entities. We remark that training tree-based models in the collaborative setting requires communication between the different members, thus increasing potential leakage compared to a non-collaborative scenario. Three types of leakage can occur:

Data Leakage. This directly endangers each party's raw training data. Data leakage can be the result of an improperly configured protection mechanism. For instance, works on random perturbation [7, 49] offer no formal privacy guarantees and are vulnerable

to reconstruction attacks [93]. As mentioned in §6.1, these approaches employ ad-hoc methods to obfuscate the training data. Similarly, encryption as obfuscation using deterministic encryption can provide only limited guarantees, as the relationship between data points is preserved. Works using DP mechanisms in the central model also inherently leak raw training data to the server; by design the server aggregates the raw data before applying the protection. Data leakage might also occur from the usage of a trained model. Shokri *et al.* propose a membership inference attack revealing if a sample was in the training set by looking at the classifier's prediction outputs [155]. Similarly, inversion attacks reveal if a sensitive attribute is used in the decision-tree [64] and property-inference attacks, if the training set satisfies some properties [13]. We consider these attacks out of scope in this analysis; they concern the inference phase and not the training.

Model Leakage. Tree-based models can be sensitive due to their business value and because they contain information about the data on which they were trained. For instance, Zhu and Du [193] quantify the privacy risks associated with publishing decision trees and show that a maximum entropy estimate can leak information about the training data. As a result, it is crucial for privacy-preserving solutions to protect the trained model. In our literature analysis, we identify works that protect the final model and its weights from any party, enabling only oblivious predictions on it. This is the case for a handful of works that use HE [5, 12, 75, 119, 123, 168], perturbation [85], and others that protect the model with DP mechanisms (see §6.2). Works publishing decision trees with DP guarantees introduce a different trade-off compared to cryptographic approaches. The privacy budget determines the amount of admissible leakage for the model and training data, at the cost of accuracy.

Leakage of Intermediate Values. As discussed in §6, numerous works tolerate the exchange of cleartext information during training to cope with challenging cryptographic operations, *e.g.*, comparisons and gains. We identify two classes depending on whether the revealed information is local (*i.e.*, related to only one party) or global (*i.e.*, aggregated over multiple parties). The type of leaked intermediate information varies depending on the learning algorithm. For standard-tree learning algorithms such as ID3, this information represents the selected split (*i.e.*, feature and split value) but also the gains and counts per tree path. For completely random approaches, the intermediate values are the class counts

at the leaves. Other statistics, such as the number of local samples reaching a specific node, residual accuracy errors [56], the “owner” of a node in the vertically distributed scenario (*i.e.*, which party holds the selected attribute for the node), might also be leaked. For gradient boosting ensemble models, the intermediate values are the weight updates or residuals. As shown in recent work by Hitaj *et al.* for deep models, the use of a DP mechanism can have a very limited effect on the privacy when such intermediate values are shared [81].

Summary. Following our observations, we systematize the different works in the literature. For each work, we first identify the privacy-preserving mechanisms employed. Then, we analyze the type of data exchanged during the learning process and under which form (*e.g.*, encrypted, blinded, aggregated, noisy, or in clear). We assess if this corresponds to leakage of (local or global) intermediate values that we classify as statistics or weights depending on the type of data and the learning algorithm. Finally, we identify which entity gains more information from the leak (*i.e.*, the beneficiary of the leak). We introduce an additional category for works that reveal the local model to the aggregator [10, 103, 180]. Finally, we describe how each work suffers from the leakage of information. Our analysis identifies three major groups: (a) works that do not leak any intermediate values, (b) works that leak intermediate values with local DP guarantees, and (c) works that leak intermediate statistics or weights. We present our findings in Table 4. For conciseness, we do not display works envisioned in the central collaborative model that use DP mechanisms or perturbation techniques for privacy preservation, as we have already analyzed their leakage above in the data leakage description. Overall, we find that only a few cryptographic works keep the final model secret [8, 12, 38, 75, 102, 119, 123, 168]. Additionally, our framework shows that, though most works of Table 4 rely on the same PETs (*e.g.*, secret sharing or HE), only a few of them avoid leaking intermediate values [5, 12, 107, 111, 121, 123, 168, 179, 181, 182]. This is probably due to the fact that ensuring *zero-leakage* during tree-based model induction comes with high communication and computation overheads.

9 Evaluation and Comparison

Overall, fair comparisons between works are non-trivial, as the learning, collaborative, and privacy models are orthogonal, and the choice of evaluation metrics and

datasets might differ. Additionally, very few works provide a public implementation which makes comparisons harder. Several datasets are commonly used for benchmarking the solutions (*e.g.*, [43, 87, 88]), but unfortunately, they often are too limited in the number of features and samples to be realistic for modern big-data tasks (see Appendix A for more details).

A formal comparison of communication and computation overhead in the literature is challenging to achieve as works seldomly vary only one of the four axes we have identified. Nonetheless, several conclusions can be drawn. The communication overhead is constrained by the collaborative model, the type of data partitioning, and the number of nodes in the tree: For instance, in the fully distributed models with N parties and horizontal data partitioning, local information needs to be exchanged with all other parties, whereas in the aggregator model, only the aggregator needs to be informed, yielding a communication in $O(N^2)$ versus $O(N)$, respectively (*c.f.* [48] vs. [12]). Conversely, in the vertical setting, communication is less affected, as the difficulty lies only in finding the best party to split the data (which can be achieved by only sharing the best locally computed gain and not intermediate values such as counts). With respect to computation overhead, the horizontal data-partitioning setting requires each party to compute local intermediates: Instead of computing this information once, it is computed locally by each party on smaller datasets. Eventually, all the intermediates are exchanged and aggregated to find the best split. Protection mechanisms also play a crucial role in the overhead. Works employing differentially private mechanisms, randomization or perturbation techniques report minimal to no computation and communication overhead as the distributed learning process operates similarly as if the data was not perturbed. Hardware-based solutions present a computational overhead ranging between $4.5\times$ and $5.1\times$ to perform encrypted XGBoost compared to the standard algorithm [101]. Cryptographic solutions have an even more dramatic impact on computational and communication overhead. For fully distributed tree induction with horizontal partitioning, SMC solutions induce a communication overhead of $O(N^2)$ (*e.g.*, [38, 100]) and while some HE solutions have a similar asymptotic overhead (*e.g.*, [149, 182]), others can reduce this overhead to $O(c^2 + Nc)$ (*e.g.*, [187]), where c is the number of class labels. However, the more complex homomorphic operations imply a bigger computational overhead, *e.g.*, an overhead of $O(N^2|A|)$ [149, 182] vs. $O(|A|)$ for a pure SMC solution [48], with $|A|$ the number of at-

Table 4. Intermediate value leakage in the literature. DP: Differential Privacy, P: Party, A: Aggregator, M: Miner.

Ref.	PET	What is leaked									Description
		Beneficiary			Perturb. values	Loc. Model	Statistics		Weights		
		P	A	M			Loc.	Glob.	Loc.	Glob.	
[104]	LSH	•								•	Table of similar instances and local gradients
[37]	LHE	•								•	Best local gain and node owner
[56]	HE+OT	•								•	P_1 learns residual errors from P_2
[70]	SMC[99]	•				•		•			Owner, number, or local best weak classifier
[71]	LHE	•						•			Best local gain
[160]	random+SMC	•						•			Best k attributes on augmented dataset to P
[8]	FHE	•							•		All gains
[9]	LHE+Blindings	•							•		All gains
[17]	SMC	•							•		All gains
[41]	SMC+SS	•							•		All gains
[48]	SSS	•							•		Depends how max computed
[95]	LHE	•							•		Conditional entropy, global class count
[96]	LHE+SS	•							•		Structure and counts of every class per leaf
[100]	SMC	•							•		Who owns best attribute
[109]	HE	•							•		Tree structure, best gain and node owner
[110]	LHE (thrs)	•							•		Best attribute
[122]	SS+GC	•							•		All or best gain
[149]	LHE+SMC	•							•		Sign of secret
[153]	SMC+SS	•							•		Global square ratio
[175]	SS		•							•	Best local gain and node owner
[119]	LHE (thrs)		•							•	Global gradients (h_i, g_i)
[166]	LHE		•							•	Local Split vector W to A
[187]	LHE+Blindings		•							•	Class distribution of each node, local best gain
[186]	LHE+Blindings		•							•	Leader learns permuted decryption
[32]	LHE		•							•	Leader learns permuted decryption
[52]	SMC		•							•	Residuals, node owner, split, and d
[72]	SMC		•							•	Node owner and index of best feature
[80]	LHE		•							•	Global gain for each feature to Leader
[113]	LHE+SS (thrs)		•							•	Gain value for all features to Leader
[154]	SMC		•							•	Gain value for all features
[156]	LHE		•							•	Gain value for all features
[163]	DP+HE		•							•	Global class distribution to Leader
[167]	LHE		•							•	Global gradients (h_i, g_i)
[118]	LHE+SS (thrs)		•							•	Best site, local best gain
[10]	Electronic Voting		•			•				•	Global gradients (h_i, g_i)
[103]	DP-GAN		•			•				•	Local ensemble
[54]	LHE		•							•	locally trained optimal tree
[53]	LHE		•							•	All gains
[190]	DP mechanism	•				•					All gains, node owner
[116]	DP mechanism		•			•					Local DP models
[180]	DP mechanism		•			•					DP protected intermediates
[162]	DP+LHE (thrs)		•			•					Local DP models
											Perturbed global gradients (more with collusion)

tributes depending only on the training set. Recently, Wu *et al.* [179] showed that combining HE and SMC can reduce the computational overhead achieving up to a $19.8\times$ speed up compared to a pure SMC solution. Indeed, in their approach, the secret-shared values among the parties are reduced from $O(n|A|)$ to $O(b|A|)$ with b the number of split values for any feature.

10 Open Challenges

Our deep analysis of privacy-preserving collaborative solutions for decision-tree induction enables us to identify some open problems and challenges in this field.

Motivating Privacy-Preserving Induction and Quantification of Leakage. As discussed in §7, the motivation for privacy-protection in collaborative tree induction is not clearly stated in the literature; there exist no papers that investigate privacy attacks that

can occur while training a tree-based model in a federated setting. On the contrary, research on other learning algorithms, *e.g.*, neural networks, demonstrates several attacks, *e.g.*, membership or property-inference ones [81, 128, 135], that might take place in collaborative settings. Indeed, the need for strong protection of intermediate computation values during tree induction has not been extensively studied. Hence, the design of more efficient protocols for privacy protection requires an understanding of the quantitative effect of such leakage on the data and model privacy. Our leakage analysis framework is a first step towards this research direction.

End-to-End Protection. Although numerous cryptographic solutions on privacy-preserving distributed decision-tree learning exist, very few of them provide full privacy protection by safeguarding the parties’ data, the training process, its outputs, and preventing intermediate value leakage. Works that achieve this protection level cover only some collaborative models and learning

algorithms. For instance, random forests and CRTs have only been covered in the offloading [12, 123] and fully-distributed [168] scenarios. The end-to-end solution by De Hoogh *et al.* works only for categorical data [38] and the work lifting this limitation operates in the offloading model [5]. Thus, hybrid solutions combining HE or secret-sharing with differential privacy to trade-off some accuracy for end-to-end privacy protection in collaborative tree-based learning is a promising direction [179].

Malicious Members in the Collective. In §7, we analyzed the literature’s threat models and found that most works consider semi-honest entities. Whereas this model is deemed acceptable for research prototypes, it might not reflect real-world collaborative scenarios involving multiple entities; some of them might have interest in exploiting the collaborative training process for their own benefits. To this end, it would be valuable to design solutions in stronger models, *e.g.*, the anytrust model [176] or a fully malicious one. While some works have addressed the former, they only consider limited collaborative models and learning algorithms, *e.g.*, [48, 179]. Addressing fully malicious models could be achieved by employing additional cryptographic techniques, *e.g.*, verifiable computation and zero-knowledge proofs, or robust learning techniques that prevent malicious participants from actively biasing the learning. For instance, zero-knowledge proofs can ensure the correct execution of decision-tree predictions on a secret model held by one party [189]. Fiore *et al.* [57, 58] apply similar techniques for ensuring correct computations on outsourced encrypted data. Alternatively, works in the direction of *Byzantine tolerance* propose statistical techniques that prevent malicious parties from poisoning the learning process [45]; these techniques could be combined with cryptographic protocols to get the best of both worlds. As noted by Bagdasaryan *et al.* [14], this is an open problem for the PETs community not only for tree-based models but for federated learning in general. We remark, however, that decision trees have a substantially different structure than neural networks and that no work has ever considered byzantine attacks for collaborative tree-based models. Consequently, understanding how byzantine parties can disrupt the learning of distributed tree-based induction is a crucial first step.

Resilience to Dropout and Fault Tolerance. Party dropouts occur when a member of the collective decides to leave the consortium while the training process has not been finalized. Following data protection laws, such as the GDPR [50], it is desirable that this

member’s data is not taken into account for the model training. This functionality, also referred to as the *right to be forgotten*, can jeopardize privacy, in particular for tree-based models [30]. Yet, not many works address this problem [118, 175]. A possible deterrent is the fact that dropouts are incompatible with widely used cryptographic solutions such as HE or secret sharing where all parties are needed to obtain a computation result. Threshold cryptosystems could offer a solution to the problem as the result of the computation can be obtained with the participation of only a subset of entities.

Fault tolerance relates to communication issues that arise when members of the collective are unavailable during the training (or respond asynchronously to queries); the unreliability of the network and potential hardware limitations of the parties could impede collaborative learning by waiting for the unavailable parties. For instance, the learning can be *stuck* if the aggregator requires a sufficient number of local updates to proceed with its computations, if the collective relies on a leader for protocol computations [187], or if a member is unavailable in the sequential setting [190]. A potential research direction is again the application of threshold cryptography and its combination with learning algorithms that cope with asynchronicity.

11 Conclusion

In this work, we have systematized the knowledge stemming from the literature on privacy-preserving collaborative tree-based model induction. We observed that preserving privacy in the distributed setting incurs challenges that most solutions address by modifying the learning algorithm or introducing new entities in the collective. Numerous privacy-preserving technologies have been employed to account for combinations of learning algorithms, collaborative settings, and threat models; hybrid solutions that combine cryptographic protocols with differential privacy are the most promising approach for achieving end-to-end privacy protection for collaborative tree model induction. We have introduced a leakage identification framework that enables protocol designers to categorize and reason about the leakage that stems from the intermediate information computed and exchanged in collaborative model training. Our systematization and framework highlighted several avenues for further research on the topic such as support for malicious models, asynchronicity, and leakage protection.

12 Acknowledgements

We would like to thank our shepherd Sebastian Meiser and the anonymous reviewers for their helpful feedback. We are also grateful to Giovanni Cherubin, Andrew Janowczyk, Peizhao Hu, and the members of the EPFL Laboratory for Data Security for their helpful comments and suggestions. This work was supported in part by the grant C17-16 (SecureKG) of the Swiss Data Science Center.

References

- [1] "Amazon sagemaker - xgboost algorithm," <https://docs.aws.amazon.com/sagemaker/latest/dg/xgboost.html>.
- [2] "DBLP: Computer science bibliography," <https://dblp.org/>.
- [3] "Google scholar," <https://scholar.google.com/>.
- [4] "Microsoft academic," <https://academic.microsoft.com/home>.
- [5] M. Abspoel, D. Escudero, and N. Volgushev, "Secure training of decision trees with continuous attributes," *Proceedings on Privacy Enhancing Technologies*, 2021.
- [6] R. Agrawal, A. Evmimievski, and R. Srikant, "Information sharing across private databases," in *ACM SIGMOD*, 2003.
- [7] R. Agrawal and R. Srikant, "Privacy-preserving data mining," in *ACM SIGMOD*, 2000.
- [8] A. Akavia, M. Leibovich, Y. S. Resheff, R. Ron, M. Shahar, and M. Vald, "Privacy-preserving decision tree training and prediction against malicious server," *iacr eprint*, 2019, <https://eprint.iacr.org/2019/1282>.
- [9] A. Alabdulkarim, M. Al-Rodhaan, T. Ma, and Y. Tian, "PPSDT: A novel privacy-preserving single decision tree algorithm for clinical decision-support systems using IoT devices," *Sensors*, 2019.
- [10] A. Alabdulkarim, M. Al-Rodhaan, Y. Tian, and A. Al-Dhelaan, "A privacy-preserving algorithm for clinical decision-support systems using random forest," *CMC Comput. Mater. Con*, 2019.
- [11] M. Andriushchenko and M. Hein, "Provably robust boosted decision stumps and trees against adversarial attacks," in *NeurIPS*, 2019.
- [12] L. J. Aslett, P. M. Esperança, and C. C. Holmes, "Encrypted statistical machine learning: new privacy preserving methods," *arXiv preprint 1508.06845*, 2015.
- [13] G. Ateniese, L. V. Mancini, A. Spognardi, A. Villani, D. Vitali, and G. Felici, "Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers," *International Journal of Security and Networks*, 2015.
- [14] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," in *AISTATS*, 2020.
- [15] R. Baghel and M. Dutta, "Privacy preserving classification by using modified c4. 5," in *IC3*. IEEE, 2013.
- [16] X. Bai, J. Yao, M. Yuan, K. Deng, X. Xie, and H. Guan, "Embedding differential privacy in decision tree algorithm with different depths," *Science China Information Sciences*, 2017.
- [17] G. Behera, "Privacy preserving c4. 5 using gini index," in *NCETACS*. IEEE, 2011.
- [18] M. Bellare, R. Dowsley, and S. Keelveedhi, "How secure is deterministic encryption?" in *PKC*, 2015.
- [19] A. Blum, C. Dwork, F. McSherry, and K. Nissim, "Practical privacy: the sulq framework," in *ACM PODS*, 2005.
- [20] M. Bojarski, A. Choromanska, K. Choromanski, and Y. Le-Cun, "Differentially-and non-differentially-private random decision trees," *arXiv preprint 1410.6973*, 2014.
- [21] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *ACM SIGSAC CCS*, 2017.
- [22] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, "(leveled) fully homomorphic encryption without bootstrapping," *ACM TOCT*, 2014.
- [23] L. Breiman, "Random forests," *Machine learning*, 2001.
- [24] L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen, *Classification and regression trees*. CRC press, 1984.
- [25] E. Bresson, D. Catalano, and D. Pointcheval, "A simple public-key cryptosystem with a double trapdoor decryption mechanism and its applications," in *ASIACRYPT*. Springer, 2003.
- [26] J. Brickell and V. Shmatikov, "Privacy-preserving classifier learning," in *FC*, 2009.
- [27] S. Bu, L. V. Lakshmanan, R. T. Ng, and G. Ramesh, "Preservation of patterns and input-output privacy," in *IEEE ICDE*, 2007.
- [28] N. Buescher, S. Boukoros, S. Bauregger, and S. Katzenbeisser, "Two is not enough: Privacy assessment of aggregation schemes in smart metering," *Proceedings on Privacy Enhancing Technologies*, 2017.
- [29] H. Chen, W. Dai, M. Kim, and Y. Song, "Efficient multi-key homomorphic encryption with packed ciphertexts with application to oblivious neural network inference," in *ACM SIGSAC CCS*, 2019.
- [30] M. Chen, Z. Zhang, T. Wang, M. Backes, M. Humbert, and Y. Zhang, "When machine unlearning jeopardizes privacy," *arXiv preprint 2005.02205*, 2020.
- [31] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *ACM SIGKDD*, 2016.
- [32] K. Cheng, T. Fan, Y. Jin, Y. Liu, T. Chen, and Q. Yang, "Secureboost: A lossless federated learning framework," *arXiv preprint 1901.08755*, 2019.
- [33] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *ASIACRYPT*. Springer, 2017.
- [34] S. Consul and S. Williamson, "Differentially private median forests for regression and classification," *arXiv 2006.08795*, 2020.
- [35] I. Damgård and M. Jurik, "A generalisation, a simplification and some applications of paillier's probabilistic public-key system," in *PKC*. Springer, 2001.
- [36] I. Damgård, V. Pastro, N. Smart, and S. Zakarias, "Multiparty computation from somewhat homomorphic encryption," in *Annual Cryptology Conference*. Springer, 2012.

- [37] J. Dansana, D. Dey, and R. Kumar, "A novel approach: Cart algorithm for vertically partitioned database in multi-party environment," in *IEEE CICT*, 2013.
- [38] S. de Hoogh, B. Schoenmakers, P. Chen, and H. op den Akker, "Practical secure decision tree learning in a tele-treatment application," in *FC*, 2014.
- [39] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," *arXiv preprint 1702.08608*, 2017.
- [40] J. Dowd, S. Xu, and W. Zhang, "Privacy-preserving decision tree mining based on random substitutions," in *ET-RICS*, 2006.
- [41] W. Du and Z. Zhan, "Building decision tree classifier on private data," in *CRPIT*, 2002.
- [42] —, "Using randomized response techniques for privacy-preserving data mining," in *ACM SIGKDD*, 2003.
- [43] D. Dua and C. Graff, "UCI machine learning repository," 2017, <http://archive.ics.uci.edu/ml>.
- [44] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*, 2006.
- [45] E.-M. El-Mhamdi, R. Guerraoui, A. Guirguis, L. N. Hoang, and S. Rouault, "Genuinely distributed byzantine machine learning," in *PODC*, 2020.
- [46] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," *CRYPTO*, 1985.
- [47] F. Emekçi, D. Agrawal, A. E. Abbadi, and A. Gulbeden, "Privacy preserving query processing using third parties," in *ICDE*. IEEE, 2006.
- [48] F. Emekçi, O. D. Sahin, D. Agrawal, and A. El Abbadi, "Privacy preserving decision tree learning over multiple parties," *Data & Knowledge Engineering*, 2007.
- [49] V. Estivill-Castro and L. Brankovic, "Data swapping: Balancing privacy against precision in mining for logic rules," in *DaWaK*. Springer, 1999.
- [50] European Parliament and Council of European Union, "General data protection regulations," 2016.
- [51] W. Fan, H. Wang, P. S. Yu, and S. Ma, "Is random model better? on its accuracy and efficiency," in *ICDM*, 2003.
- [52] W. Fang and B. Yang, "Privacy preserving decision tree learning over vertically partitioned data," in *CSSE*. IEEE, 2008.
- [53] W. Fang, B. Yang, and D. Song, "Preserving private knowledge in decision tree learning," *JCP*, 2010.
- [54] W. Fang, B. Yang, D. Song, and Z. Tang, "A new scheme on privacy-preserving distributed decision-tree mining," in *ETCS*. IEEE, 2009.
- [55] W. Fang, C. Chen, J. Tan, C. Yu, Y. Lu, L. Wang, L. Wang, J. Zhou *et al.*, "A hybrid-domain framework for secure gradient tree boosting," *arXiv preprint 2005.08479*, 2020.
- [56] Z. Feng, H. Xiong, C. Song, S. Yang, B. Zhao, L. Wang, Z. Chen, S. Yang, L. Liu, and J. Huan, "Securegbm: Secure multi-party gradient boosting," in *Big Data*. IEEE, 2019.
- [57] D. Fiore, R. Gennaro, and V. Pastro, "Efficiently verifiable computation on encrypted data," in *ACM SIGSAC CCS*, 2014.
- [58] D. Fiore, A. Nitulescu, and D. Pointcheval, "Boosting verifiable computation on encrypted data," in *PKC*, 2020, pp. 124–154.
- [59] S. Fletcher and M. Z. Islam, "A differentially private random decision forest using reliable signal-to-noise ratios," in *Australasian joint conference on artificial intelligence*, 2015.
- [60] —, "Differentially private random decision forests using smooth sensitivity," *Expert Systems with Applications*, 2017.
- [61] —, "Decision tree classification with differential privacy: A survey," *ACM CSUR*, 2019.
- [62] —, "A differentially private decision forest," in *AusDM*. CRPIT, 2015.
- [63] P. K. Fong and J. H. Weber-Jahnke, "Privacy preserving decision tree learning using unrealized data sets," *TKDE*, 2012.
- [64] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *ACM SIGSAC CCS*, 2015.
- [65] M. J. Freedman, K. Nissim, and B. Pinkas, "Efficient private matching and set intersection," in *EUROCRYPT*, 2004.
- [66] A. A. Freitas, "Comprehensible classification models: a position paper," *ACM SIGKDD explorations newsletter*, 2014.
- [67] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *J. Comput. Syst. Sci*. Springer, 1995.
- [68] A. Friedman and A. Schuster, "Data mining with differential privacy," in *ACM SIGKDD*, 2010.
- [69] J. H. Friedman, "Greedy function approximation: a gradient boosting machine," *Annals of statistics*, 2001.
- [70] S. Gambs, B. Kégl, and E. Aïmeur, "Privacy-preserving boosting," *Data Mining and Knowledge Discovery*, 2007.
- [71] A. Gangrade and R. Patel, "Building privacy-preserving c4.5 decision tree classifier on multi-parties," *IJCSE*, 2009.
- [72] —, "A novel protocol for privacy preserving decision tree over horizontally partitioned data," *IJARCS*, 2011.
- [73] —, "Privacy preserving two-layer decision tree classifier for multiparty databases," *IJCIT*, 2012.
- [74] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Machine learning*, 2006.
- [75] I. Giacomelli, S. Jha, R. Kleiman, D. Page, and K. Yoon, "Privacy-preserving collaborative prediction using random forests," *AMIA Jt Summits Transl Sci Proc.*, 2019.
- [76] C. Giannella, K. Liu, T. Olsen, and H. Kargupta, "Communication efficient construction of decision trees over heterogeneously distributed data," in *ICDM*. IEEE, 2004.
- [77] O. Goldreich, *General Cryptographic Protocols*. Cambridge University Press, 2004, vol. 2, p. 599–764.
- [78] —, *Foundations of cryptography: volume 2, basic applications*. Cambridge university press, 2009.
- [79] Z. Guan, X. Sun, L. Shi, L. Wu, and X. Du, "A differentially private greedy decision forest classification algorithm with high utility," *Computers & Security*, 2020.
- [80] S. Han and W. K. Ng, "Multi-party privacy-preserving decision trees for arbitrarily partitioned data," *Int J Intell Control Syst*, 2007.
- [81] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep models under the gan: information leakage from collaborative deep learning," in *ACM SIGSAC CCS*, 2017.
- [82] N. Homer, S. Szelinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan,

- S. F. Nelson, and D. W. Craig, "Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays," *PLoS Genet*, 2008.
- [83] J. Hou, Q. Li, S. Meng, Z. Ni, Y. Chen, and Y. Liu, "Dprf: A differential privacy protection random forest," *IEEE Access*, 2019.
- [84] G. Jagannathan, K. Pillaipakkamnatt, and R. N. Wright, "A practical differentially private random decision tree classifier," in *IEEE ICDMW*, 2009.
- [85] M. A. Kadampur *et al.*, "A noise addition scheme in decision tree for privacy preserving data mining," *arXiv preprint 1001.3504*, 2010.
- [86] Kaggle, "Credit card fraud," <https://www.kaggle.com/mlg-ulb/creditcardfraud>.
- [87] —, "Default of credit card data," <https://www.kaggle.com/uciml/default-of-credit-card-clients-dataset>.
- [88] —, "Give me some credit," <https://www.kaggle.com/c/GiveMeSomeCredit/data>.
- [89] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, "Advances and open problems in federated learning," *arXiv preprint 1912.04977*, 2019.
- [90] G. Kalyani, M. C. S. Rao, and B. Janakiramaiah, "Decision tree based data reconstruction for privacy preserving classification rule mining," *Informatica*, 2017.
- [91] M. Kantarcioglu and C. Clifton, "Privacy-preserving distributed mining of association rules on horizontally partitioned data," *IEEE TKDE*, 2004.
- [92] D. Kaplan, J. Powell, and T. Woller, "AMD memory encryption," Advanced Micro Devices, Tech. Rep., 2016.
- [93] H. Kargupta, S. Datta, Q. Wang, and K. Sivakumar, "On the privacy preserving properties of random data perturbation techniques," in *IEEE ICDM*, 2003.
- [94] M. Keller, "MP-SPDZ: A versatile framework for multi-party computation," Cryptology ePrint Archive, Report 2020/521, 2020, <https://eprint.iacr.org/2020/521>.
- [95] F. Khodaparast, M. Sheikhalishahi, H. Haghighi, and F. Martinelli, "Privacy preserving random decision tree classification over horizontally and vertically partitioned data," in *DASC/PiCom/DataCom/CyberSciTech*. IEEE, 2018.
- [96] H. Kikuchi, K. Ito, M. Ushida, H. Tsuda, and Y. Yamaoka, "Privacy-preserving distributed decision tree learning with boolean class attributes," in *AINA*. IEEE, 2013.
- [97] E. Kim, J. Jeong, H. Yoon, Y. Kim, J. Cho, and J. H. Cheon, "How to securely collaborate on data: Decentralized threshold he and secure key update," *IEEE Access*, 2020.
- [98] Á. Kiss, M. Naderpour, J. Liu, N. Asokan, and T. Schneider, "Sok: Modular and efficient private decision tree evaluation," *Proceedings on Privacy Enhancing Technologies*, 2019.
- [99] L. Kissner and D. Song, "Privacy-preserving set operations," in *Annual International Cryptology Conference*. Springer, 2005.
- [100] B. Kuijpers, V. Lemmens, B. Moelans, and K. Tuyls, "Privacy preserving ID3 over horizontally, vertically and grid partitioned data," *arXiv preprint 0803.1555*, 2008.
- [101] A. Law, C. Leung, R. Poddar, R. A. Popa, C. Shi, O. Sima, C. Yu, X. Zhang, and W. Zheng, "Secure collaborative training and inference for xgboost," *arXiv preprint 2010.02524*, 2020.
- [102] C. Leung, A. Law, and O. Sima, "Towards privacy-preserving collaborative gradient boosted decision trees," UC Berkeley, Tech. Rep., 2019.
- [103] J. Li, Y. Tian, Y. Zhu, T. Zhou, J. Li, K. Ding, and J. Li, "A multicenter random forest model for effective prognosis prediction in collaborative clinical research network," *Artificial Intelligence in Medicine*, 2020.
- [104] Q. Li, Z. Wen, and B. He, "Practical federated gradient boosting decision trees," *arXiv preprint 1911.04206*, 2019.
- [105] Q. Li, Z. Wu, Z. Wen, and B. He, "Privacy-preserving gradient boosting decision trees," in *AAAI*, 2020.
- [106] Y. Li, C. Bai, and C. K. Reddy, "A distributed ensemble approach for mining healthcare data under privacy constraints," *Information sciences*, 2016.
- [107] Y. Li, Z. L. Jiang, X. Wang, and S.-M. Yiu, "Privacy-preserving ID3 data mining over encrypted data in outsourced environments with multiple keys," in *IEEE Int. Conf. Comp. Sci. & Eng. (CSE) and IEEE EUC*. IEEE, 2017.
- [108] Y. Li, Z. L. Jiang, X. Wang, S.-M. Yiu, and J. Fang, "Outsourced privacy-preserving random decision tree algorithm under multiple parties for sensor-cloud integration," in *ISPEC*. Springer, 2017.
- [109] Y. Li, Z. L. Jiang, X. Wang, S.-M. Yiu, and P. Zhang, "Outsourcing privacy preserving ID3 decision tree algorithm over encrypted data-sets for two-parties," in *2017 IEEE Trustcom/BigDataSE/ICSS*. IEEE, 2017.
- [110] Y. Li, Z. L. Jiang, L. Yao, X. Wang, S.-M. Yiu, and Z. Huang, "Outsourced privacy-preserving c4.5 decision tree algorithm over horizontally and vertically partitioned dataset among multiple parties," *Cluster Computing*, 2019.
- [111] Y. Lindell and B. Pinkas, "Privacy preserving data mining," in *Annual International Cryptology Conference*, 2000.
- [112] L. Liu, M. Kantarcioglu, and B. Thuraisingham, "Privacy preserving decision tree mining from perturbed data," in *HICSS*. IEEE, 2009.
- [113] L. Liu, R. Chen, X. Liu, J. Su, and L. Qiao, "Towards practical privacy-preserving decision tree training and evaluation in the cloud," *IEEE TIFS*, 2020.
- [114] X. Liu, Q. Li, T. Li, and D. Chen, "Differentially private classification with decision tree ensemble," *Applied Soft Computing*, 2018.
- [115] X. Liu, R. H. Deng, K.-K. R. Choo, and J. Weng, "An efficient privacy-preserving outsourced calculation toolkit with multiple keys," *IEEE TIFS*, 2016.
- [116] Y. Liu, M. Chen, W. Zhang, J. Zhang, and Y. Zheng, "Federated extra-trees with privacy preserving," *arXiv*, 2020.
- [117] Y. Liu, Y. Liu, Z. Liu, Y. Liang, C. Meng, J. Zhang, and Y. Zheng, "Federated forest," *IEEE Trans. on Big Data*, 2020.
- [118] Y. Liu, Z. Ma, X. Liu, S. Ma, S. Nepal, and R. Deng, "Boosting privately: Privacy-preserving federated extreme boosting for mobile crowdsensing," *arXiv preprint 1907.10218*, 2019.

- [119] Y. Liu, Z. Ma, X. Liu, Z. Wang, S. Ma, and K. Ren, "Revocable federated learning: A benchmark of federated forest," *arXiv preprint 1911.03242*, 2019.
- [120] Y.-h. Liu, B.-R. Yang, D.-y. Cao, and N. Ma, "State-of-the-art in distributed privacy preserving data mining," in *IEEE ICCSN*, 2011.
- [121] P. Lory, "Enhancing the efficiency in privacy preserving learning of decision trees in partitioned databases," in *PSD*, 2012.
- [122] Q. Ma and P. Deng, "Secure multi-party protocols for privacy preserving data mining," in *WASA*, 2008.
- [123] Z. Ma, J. Ma, Y. Miao, and X. Liu, "Privacy-preserving and high-accurate outsourced disease predictor on random forest," *Information Sciences*, 2019.
- [124] S. Mabu, M. Obayashi, and T. Kuremoto, "Ensemble learning of rule-based evolutionary algorithm using multi-layer perceptron for supporting decisions in stock trading problems," *Applied soft computing*, 2015.
- [125] F. McKeen, I. Alexandrovich, A. Berenzon, C. V. Rozas, H. Shafi, V. Shanbhogue, and U. R. Savagaonkar, "Innovative instructions and software model for isolated execution." *Hasp@ isca*, 2013.
- [126] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *AISTATS*, 2017.
- [127] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *FOCS*. IEEE, 2007.
- [128] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting unintended feature leakage in collaborative learning," in *IEEE S&P*, 2019.
- [129] R. Mitchell, "Gradient boosting, decision trees and xgboost with cuda," <https://developer.nvidia.com/blog/gradient-boosting-decision-trees-xgboost-cuda/>.
- [130] N. Mohammed, R. Chen, B. C. Fung, and P. S. Yu, "Differentially private data release for data mining," in *ACM SIGKDD*, 2011.
- [131] C. Molnar, *Interpretable Machine Learning*. Lulu. com, 2020.
- [132] C. Mouchet, J. Troncoso-Pastoriza, and J.-P. Hubaux, "Multiparty homomorphic encryption: From theory to practice," *eprint*, 2020, <https://eprint.iacr.org/2020/304>.
- [133] M. Naor and B. Pinkas, "Oblivious transfer and polynomial evaluation," in *STOC*, 1999.
- [134] —, "Oblivious polynomial evaluation," *SIAM Journal on Computing*, 2006.
- [135] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," in *IEEE S&P*, 2019.
- [136] M. Naveed, S. Kamara, and C. V. Wright, "Inference attacks on property-preserving encrypted databases," in *ACM SIGSAC CCS*, 2015.
- [137] R. Nock and W. Henecka, "Boosted and differentially private ensembles of decision trees," *arXiv preprint 2001.09384*, 2020.
- [138] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *EUROCRYPT*. Springer, 1999.
- [139] N. Papernot, P. McDaniel, A. Sinha, and M. P. Wellman, "Sok: Security and privacy in machine learning," in *Eu-roS&P*. IEEE, 2018.
- [140] A. Patil and S. Singh, "Differential private random forest," in *ICACCI*. IEEE, 2014.
- [141] G. Piatetsky, "Top data science and machine learning methods used in 2017," <https://www.kdnuggets.com/2017/12/top-data-science-machine-learning-methods.html>.
- [142] B. Pinkas, T. Schneider, and M. Zohner, "Faster private set intersection based on OT extension," in *USENIX Security Symposium*, 2014.
- [143] A. Pyrgelis, C. Troncoso, and E. De Cristofaro, "What does the crowd say about you? evaluating aggregation-based location privacy," *Proceedings on privacy enhancing technologies*, 2017.
- [144] J. R. Quinlan, "Induction of decision trees," *Machine learning*, 1986.
- [145] J. Quinlan, *C4. 5: programs for machine learning*. Elsevier, 1993.
- [146] S. Rana, S. K. Gupta, and S. Venkatesh, "Differentially private random forest with high utility," in *ICDM*, 2015.
- [147] G. N. Rao, M. S. Harini, and C. R. Kishore, "A cryptographic privacy preserving approach over classification," in *ICT and Critical Infrastructure: Proceedings of the 48th Annual Convention of Computer Society of India-Vol II*. Springer, 2014.
- [148] A. Salem, A. Bhattacharya, M. Backes, M. Fritz, and Y. Zhang, "Updates-leak: Data set inference and reconstruction attacks in online learning," in *USENIX Security Symposium*, 2020.
- [149] S. Samet and A. Miri, "Privacy preserving ID3 using gini index over horizontally partitioned data," in *AICCSA*, 2008.
- [150] V. Sandulescu and M. Chiru, "Predicting the future relevance of research institutions-the winning solution of the KDD cup 2016," *arXiv preprint 1609.02728*, 2016.
- [151] A. Shamir, "How to share a secret," *Communications of the ACM*, 1979.
- [152] S. Sharma and A. S. Rajawat, "A secure privacy preservation model for vertically partitioned distributed data," in *ICTBIG*. IEEE, 2016.
- [153] M. A. Sheela and K. Vijayalakshmi, "A novel privacy preserving decision tree induction," in *IEEE CICT*, 2013.
- [154] Y. Shen, H. Shao, and L. Yang, "Privacy preserving C4.5 algorithm over vertically distributed datasets," in *NSWCTC*. IEEE, 2009.
- [155] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *IEEE S&P*, 2017.
- [156] L. Sumalatha and P. U. Sankar, "Fuzzy random decision tree (frdt) framework for privacy preserving data mining," in *SAI*. IEEE, 2016.
- [157] E. Suthampan and S. Maneewongvatana, "Privacy preserving decision tree in multi party environment," in *AIRS*, 2005.
- [158] K. Tandel and J. N. Patel, "Privacy preserving decision tree classification on horizontal partition data," *IJERT*, 2016.
- [159] V. D. Team, "VIFF, the virtual ideal functionality framework."
- [160] Z. Teng and W. Du, "A hybrid multi-group privacy-preserving approach for building decision trees," in *PAKDD*, 2007.

- [161] The Health Insurance Portability and Accountability Act of 1996, 2014, www.hhs.gov/hipaa.
- [162] Z. Tian, R. Zhang, X. Hou, J. Liu, and K. Ren, "Federboost: Private federated learning for gbdtd," *arXiv preprint 2011.02796*, 2020.
- [163] S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, R. Zhang, and Y. Zhou, "A hybrid approach to privacy-preserving federated learning," in *ACM AISec*, 2019.
- [164] S. Truex, L. Liu, M. E. Gursoy, and L. Yu, "Privacy-preserving inductive learning with decision trees," in *Big-Data Congress*, 2017.
- [165] P. S. Vadivu and S. Nithya, "An improved privacy preserving with rsa and c5.0 decision tree learning for unrealized datasets," *International Journal*, 2014.
- [166] J. Vaidya and C. Clifton, "Privacy-preserving decision trees over vertically partitioned data," in *IFIP DBSec*. Springer, 2005.
- [167] J. Vaidya, C. Clifton, M. Kantarcioglu, and A. S. Patterson, "Privacy-preserving decision trees over vertically partitioned data," *ACM TKDD*, 2008.
- [168] J. Vaidya, B. Shafiq, W. Fan, D. Mehmood, and D. Lorenzi, "A random decision tree framework for privacy-preserving data mining," *TDSC*, 2013.
- [169] J. Van Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wensch, Y. Yarom, and R. Strackx, "Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution," in *USENIX Security Symposium*, 2018.
- [170] V. S. Verykios, E. Bertino, I. N. Fovino, L. P. Provenza, Y. Saygin, and Y. Theodoridis, "State-of-the-art in privacy preserving data mining," *ACM Sigmod Record*, 2004.
- [171] K. W. Wang, T. Dick, and M.-F. Balcan, "Scalable and provably accurate algorithms for differentially private distributed decision tree learning," *AAAI PPAI*, 2020.
- [172] K. Wang, Y. Xu, R. She, and P. S. Yu, "Classification spanning private databases," in *AAAI*, 2006.
- [173] S. Wang and J. M. Chang, "Privacy-preserving boosting in the local setting," *arXiv preprint 2002.02096*, 2020.
- [174] W. Wang, G. Chen, X. Pan, Y. Zhang, X. Wang, V. Bind-schaedler, H. Tang, and C. A. Gunter, "Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX," in *ACM SIGSAC CCS*, 2017.
- [175] Z. Wang, Y. Yang, Y. Liu, X. Liu, B. B. Gupta, and J. Ma, "Cloud-based federated boosting for mobile crowdsensing," *arXiv preprint 2005.05304*, 2020.
- [176] D. I. Wolinsky, H. Corrigan-Gibbs, B. Ford, and A. Johnson, "Scalable anonymous group communication in the anytrust model," Naval Research Lab Washington DC, Tech. Rep., 2012.
- [177] D. Wu, T. Wu, and X. Wu, "A differentially private random decision tree classifier with high utility," in *ML4CS*. Springer, 2020.
- [178] D. J. Wu, T. Feng, M. Naehrig, and K. Lauter, "Privately evaluating decision trees and random forests," *Proceedings on Privacy Enhancing Technologies*, 2016.
- [179] Y. Wu, S. Cai, X. Xiao, G. Chen, and B. C. Ooi, "Privacy preserving vertical federated learning for tree-based models," *VLDB*, no. 11, 2020.
- [180] T. Xiang, Y. Li, X. Li, S. Zhong, and S. Yu, "Collaborative ensemble learning under differential privacy," in *Web Intelligence*, 2018.
- [181] M.-J. Xiao, K. Han, L.-S. Huang, and J.-Y. Li, "Privacy preserving c4. 5 algorithm over horizontally partitioned data," in *GCC*. IEEE, 2006.
- [182] M.-J. Xiao, L.-S. Huang, Y.-L. Luo, and H. Shen, "Privacy preserving ID3 algorithm over horizontally partitioned data," in *PDCAT*. IEEE, 2005.
- [183] B. Xin, W. Yang, S. Wang, and L. Huang, "Differentially private greedy decision forest," in *ICASSP*. IEEE, 2019.
- [184] M. Yang, L. Song, J. Xu, C. Li, and G. Tan, "The tradeoff between privacy and accuracy in anomaly detection using federated xgboost," *arXiv preprint 1907.07157*, 2019.
- [185] A. C.-C. Yao, "How to generate and exchange secrets," in *SFCS*. IEEE, 1986.
- [186] J. Zhan, "Using homomorphic encryption for privacy-preserving collaborative decision tree classification," in *IEEE CIDM*, 2007.
- [187] J. Zhan, S. Matwin, and L. Chang, "Privacy preserving decision tree classification over horizontally partitioned data," in *ICEB*, 2005.
- [188] C. Zhang, Y. Li, and Z. Chen, "Dpsets: A differentially private extratrees," in *CIS*, 2017.
- [189] J. Zhang, Z. Fang, Y. Zhang, and D. Song, "Zero knowledge proofs for decision tree predictions and accuracy," in *ACM SIGSAC CCS*, 2020.
- [190] L. Zhao, L. Ni, S. Hu, Y. Chen, P. Zhou, F. Xiao, and L. Wu, "Inprivate digging: Enabling tree-based distributed data mining with differential privacy," in *IEEE INFOCOM*, 2018.
- [191] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," in *NeurIPS*, 2019.
- [192] T. Zhu, P. Xiong, Y. Xiang, and W. Zhou, "An effective differentially private data releasing algorithm for decision tree," in *IEEE TrustCom*, 2013.
- [193] Z. Zhu and W. Du, "Understanding privacy risk of publishing decision trees," in *IFIP DBSec*. Springer, 2010.

A Experimental Evaluations

We now expand our analysis with respect to the evaluations performed by the surveyed works.

Model Evaluation Metrics. Most works evaluate their solutions against the model's accuracy, error, area under the curve or F1 score on standard open source datasets [43, 86–88]. Most DP-based works estimate the model's accuracy as a function of the privacy budget enabling cross-work comparison [84, 140, 188]. Another popular metric for cryptographic solutions is the training execution time that is evaluated against the number of trees, the depth, or the number of parties.

Datasets. The UCI database [43] and Kaggle [86–88], are the predominant sources of the datasets used in the literature. Although this facilitates comparisons, we

note that several of these datasets contain a very limited number of features or very few samples, hence they might not be representative of modern big-data tasks performed today.

Open Source Implementations. Finally, very few works provide a public implementation of their solution [20, 34, 59, 60, 62, 102, 184]; this makes reproduction and comparison with related works harder.

B Literature Analysis

In this section, we provide further details about the different works studied in this paper. As discussed in §2, we used Google Scholar [3], Microsoft Academic [4], and DBLP [2] to discover works related to combinations of the different learning algorithms with the following key words: $\{private, privacy-preserving, collaborative, distributed, training, induction, learning\}$. Then, we cross-referenced the different works to expand the search. Following our search methodology we discovered 103 works originating from a wide variety of communities. Figure 1 shows the distribution of the investigated works across their communities. We find that most of the works come from the *machine learning*, and *data mining* communities (ML in Fig. 1): These account for 36% of the investigated works. The *security, privacy, and cryptography* community (SP) hosts 23% of the works with 24 publications. General *computer science* journals are also well represented with 16 works (dubbed CS in Fig. 1). *Systems* cover a wide range of fields from distributed to intelligent systems and represent 10% of the literature on collaborative privacy-preserving tree-based model learning. The *communications, medical, and database* communities account respectively for 6%, 2%, and 2% of the surveyed works. We regroup the remaining works as *others* (e.g., ubiquitous computing, signal processing, etc.). At the time of writing, the ML and SP communities appear as the most influential ones on the topic, yielding more than 94% of the citations across all communities.

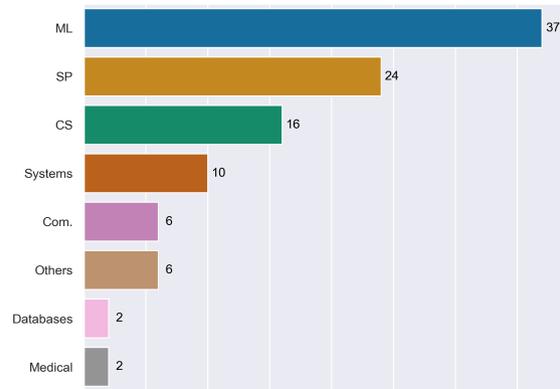


Fig. 1. Number of surveyed works across the different communities.