
Accuracy, Interpretability, and Differential Privacy via Explainable Boosting

Harsha Nori¹ Rich Caruana¹ Zhiqi Bu² Judy Hanwen Shen³ Janardhan Kulkarni¹

Abstract

We show that adding differential privacy to Explainable Boosting Machines (EBMs), a recent method for training interpretable ML models, yields state-of-the-art accuracy while protecting privacy. Our experiments on multiple classification and regression datasets show that DP-EBM models suffer surprisingly little accuracy loss even with strong differential privacy guarantees. In addition to high accuracy, two other benefits of applying DP to EBMs are: a) trained models provide exact global and local interpretability, which is often important in settings where differential privacy is needed; and b) the models can be edited after training without loss of privacy to correct errors which DP noise may have introduced.

1. Introduction

Security researchers have repeatedly shown that machine learning models can leak information about training data (Carlini et al., 2018; Melis et al., 2019). In industries like healthcare, finance and criminal justice, models are trained on sensitive information, and this form of leakage can be especially disastrous. To combat this, researchers have embraced differential privacy, which establishes a strong mathematical standard for privacy guarantees on algorithms (Dwork et al., 2006; 2014). In many of these high-stakes situations, model interpretability is also important to provide audits, help domain experts such as doctors vet the models, and to correct unwanted errors before deployment (Caruana et al., 2015; Rudin, 2019). In this paper, we address both concerns by developing a private algorithm for learning Generalized Additive Models (GAMs) (Hastie & Tibshirani, 1990). We show that this method can provide strong privacy guarantees, high accuracy, and exact global and local interpretability on tabular datasets.

¹Microsoft, Redmond, USA. ²University of Pennsylvania, Philadelphia, USA. ³Stanford University, Palo Alto, USA. Correspondence to: Harsha Nori <hanori@microsoft.com>.

While GAMs were traditionally fit using smooth low-order splines (Hastie & Tibshirani, 1990), we focus on Explainable Boosting Machines (EBMs), a modern implementation that learns shape functions using boosted decision trees (Lou et al., 2012; Nori et al., 2019). EBMs are especially interesting because they often match the accuracy of complex blackbox algorithms like XGBoost and random forests, while having a simple optimization procedure and final structure (Chang et al., 2020; Wang et al., 2020).

Our main contributions for this paper are:

- We introduce DP-EBMs, a differentially private version of EBMs, and provide a rigorous privacy analysis of this algorithm using the recently introduced GDP framework (Dong et al., 2019).
- Our experimental results on tabular classification and regression problems show that DP-EBMs significantly outperform other DP learning methods. For example, at $\epsilon = 0.5$, DP-EBMs have at most a 0.05 loss in AUC compared to non-private EBMs on benchmark datasets.
- We demonstrate how combining interpretability with differential privacy can address common concerns with DP in practice by enabling users to repair some of the impact of noise on the model and enforce desirable constraints like monotonicity.

Before diving into details in the following sections, Figure 1 provides a quick peek at the empirical results. In summary, DP-EBMs outperform other differentially private learning methods on a variety of classification and regression tasks for many reasonable values of ϵ .

2. Preliminaries

2.1. Explainable Boosting Machines

Explainable Boosting Machines belong to the family of Generalized Additive Models (GAMs), which are restricted machine learning models that have the form:

$$g(E[y]) = \beta + f_0(x_0) + f_1(x_1) + \dots + f_k(x_k)$$

where β is an intercept, each f_j is a univariate function that operates on a single input feature x_j , and g is a link function that adapts the model to different settings like classification and regression (Hastie & Tibshirani, 1990).




Figure 1. Comparison of two variants of DP-EBMs with other DP algorithms on four classification datasets. DP-EBMs significantly outperform other methods in every setting. For the full experimental setup and comparisons on regression models, please see Section 4.

While GAMs are more flexible than linear models (where each function f_j is further restricted to be linear), they are significantly less flexible than most machine learning models due to their inability to learn high-order interactions between features (e.g. $f(x_0, x_1, x_2)$). This restricted additive structure has the benefit of allowing GAMs to provide exact interpretability. At prediction time, each feature contributes a score, which are then summed and passed through a link function. These scores show exactly what each feature contributes to the prediction, and can be sorted, compared, and reasoned about (Lundberg & Lee, 2017). In addition, each function f_k can be visualized to provide an exact global description of how the model operates across varying inputs.

EBMs are a recent, popular open-source implementation of boosted-tree GAMs (Nori et al., 2019). We extend the EBM package to include DP-EBMs¹, which makes DP-EBMs as easy to use as regular EBMs or any scikit-learn model.

The EBM training procedure begins by bucketing data from continuous features into discrete bins, ensuring that each bin has approximately equal amounts of data. This pre-processing step is a common optimization in tree-based learning algorithms, and is used by popular packages like LightGBM and XGBoost (Ke et al., 2017; Chen & Guestrin, 2016). The most time consuming part of training a decision tree is finding the best split; discretizing the data before growing trees reduces the search space for splits which can significantly speed up learning with little cost in accuracy.

After pre-processing, the goal is to learn shape functions f_k for each feature. In traditional boosting, each tree greedily searches the feature space for the next best feature to split. In contrast, EBMs use *cyclic* gradient boosting to visit each

¹<https://github.com/interpretml/interpret>

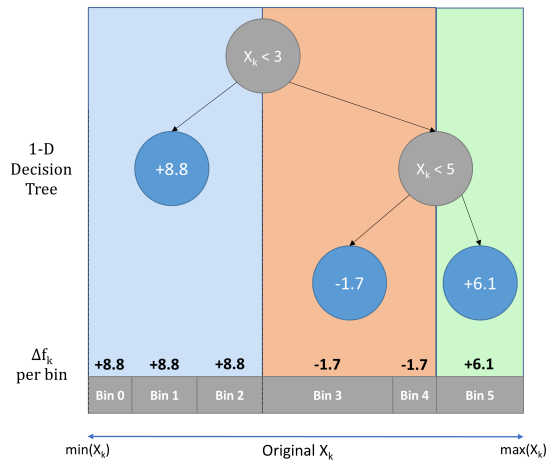


Figure 2. A single iteration of cyclic boosting, showing how each tree operates on pre-processed data and gets collapsed into the univariate shape function f_k .

feature in round-robin fashion. To enforce additivity, each tree is only allowed to use one feature, thus preventing interactions from being learned (Lou et al., 2012).

Cyclic boosting begins by growing a shallow decision tree on the first feature in the dataset. The predictions the tree makes on each bin of the histogram are then multiplied by a low learning rate, and these weak predictions for each bin of data become the initial shape function for the first feature. The process then iterates to the second feature, where a tree is trained to predict on the residuals (the remaining error) of the first feature’s model. Once a shallow tree has been learned for every feature, the boosting process cycles back to the first feature and continues in a round robin fashion for all E epochs to jointly optimize all functions. The pseudocode for this algorithm can be found in Algorithm 1.

Algorithm 1 Explainable Boosting

```

1: Input: data  $X$ , labels  $y$ , epochs  $E$ ,
   learning rate  $\eta$ , max splits  $m$ 
2: Output: 1d functions  $f_k$  per feature
3:
4:  $t = 0$ 
5: Initialize residuals:  $r_i^t = y_i$ 
6: for feature  $0 \dots K$  do
7:   Bin data:  $H_k = \text{Bin}(X[:, k])$ 
8:   Initialize output function:  $f_k^t = [0, \dots, 0]$ 
9: end for
10:
11: for epoch  $1 \dots, E$  do
12:   for feature  $0, \dots, K$  do
13:      $t += 1$ 
14:     Select best splits  $S_0, \dots, S_m$ 
15:     for split  $\ell \in \{0, \dots, m\}$  do
16:       Sum residuals:  $T = \eta \cdot \sum_{b \in S_\ell} \sum_{x_i \in H_k(b)} r_i^t$ 
17:       ...
18:       Calculate average:  $\mu = \frac{T}{\sum_{b \in S_\ell} H_k(b)}$ 
19:       for each histogram bin  $b \in S_\ell$  do
20:         Update output function:  $f_k^t(b) = f_k^t(b) + \mu$ 
21:         ...
22:       end for
23:     end for
24:     for each data point  $x_i$  do
25:       Residuals:  $r_i^{t+1} = y_i - \sum_k f_k^t(\rho(H_k, x_i))$ 
26:     end for
27:   end for
28: end for

```

2.2. Differential Privacy

Here we state some basic results in Differential Privacy (DP) that we use in our analysis.

Definition 1 (Differential Privacy). *A randomized algorithm \mathcal{A} is (ϵ, δ) -differentially private if for all neighboring databases $D_1, D_2 \in D^n$, and for all sets S of possible outputs:*

$$\Pr[\mathcal{A}(D_1) \in S] \leq e^\epsilon \Pr[\mathcal{A}(D_2) \in S] + \delta \quad (1)$$

Theorem 1 (Gaussian Mechanism (Dwork et al., 2014)). *Given any function $f : D \rightarrow \mathbb{R}^k$, the Gaussian Mechanism is defined as:*

$$\mathcal{M}(x, f(\cdot), \epsilon, \delta) = f(x) + (Y_1, \dots, Y_k) \quad (2)$$

where Δ_2 is the ℓ_2 -sensitivity and Y_i are i.i.d. random variables drawn from $\mathcal{N}(0, \sigma^2)$. The Gaussian Mechanism is (ϵ, δ) -differentially private when $\sigma > \sqrt{2 \ln 1.25 / \delta} \Delta_2 / \epsilon$ and $\epsilon \in (0, 1)$.

Algorithm 2 Differentially Private Explainable Boosting

```

1: Input: data  $X$ , labels  $y$ , epochs  $E$ , learning rate  $\eta$ , max
   splits  $m$ , range of labels  $R$ , privacy parameters  $\epsilon, \delta$ 
2: Output: 1d functions  $f_k$  per feature
3:
4:  $t = 0$ 
5: Initialize residuals:  $r_i^t = y_i$ 
6: for feature  $0 \dots K$  do
7:   Privately bin data:  $\hat{H}_k = \text{DPBin}(X[:, k], \epsilon_{bin})$ 
8:   Initialize output function:  $f_k^t = [0, \dots, 0]$ 
9: end for
10:
11: for epoch  $1 \dots, E$  do
12:   for feature  $0, \dots, K$  do
13:      $t += 1$ 
14:     Randomly select splits  $S_0, \dots, S_m$ 
15:     for split  $\ell \in \{0, \dots, m\}$  do
16:       Sum residuals:  $T = \eta \cdot \sum_{b \in S_\ell} \sum_{x_i \in \hat{H}_k(b)} r_i^t$ 
17:       Add noise:  $\hat{T} = T + \sigma \cdot \eta R \cdot \mathcal{N}(0, 1)$ 
18:       Calculate private average:  $\mu = \frac{\hat{T}}{\sum_{b \in S_\ell} \hat{H}_k(b)}$ 
19:       for each histogram bin  $b \in S_\ell$  do
20:         Update output function:  $f_k^t(b) = f_k^t(b) + \mu$ 
21:         We release  $f_k^t(b)$  values publicly.
22:       end for
23:     end for
24:     for each data point  $x_i$  do
25:       Residuals:  $r_i^{t+1} = y_i - \sum_k f_k^t(\rho(\hat{H}_k, x_i))$ 
26:     end for
27:   end for
28: end for

```

One of the main strengths of DP are the composition theorems, which analyze the cumulative privacy guarantee when applying many differentially private mechanisms.

Theorem 2. (Theorem 4.3 from (Kairouz et al., 2017))

For $\Delta > 0$, $\epsilon > 0$ and $\delta \in [0, 1]$, the mechanism that adds Gaussian noise with variance:

$$8k\Delta^2 \log(e + (\epsilon/\delta)) / \epsilon^2 \quad (3)$$

satisfies (ϵ, δ) -differential privacy under k -fold adaptive composition.

A qualitative way to understand the above theorem is that if there are k differentially private mechanisms each of which is (ϵ, δ) -DP acting on the same data set, then the overall privacy loss is roughly $\epsilon \cdot \sqrt{k}$.

Unfortunately, Theorem 2 is not the optimal bound one can achieve on the composition of private mechanisms. A tighter analysis of composition for Gaussian mechanisms, called Gaussian Differential Privacy (GDP), was recently proposed by (Dong et al., 2019). In our experiments, GDP analysis gave us better privacy bounds. We summarize

the main theorems we borrow from (Dong et al., 2019) below. For completeness, we compare the results from both composition methods (“EBM-Classic” and “EBM-GDP”) in Tables 2 and 3.

Theorem 3. *For a dataset D , define the Gaussian mechanism that operates on a univariate statistic θ with sensitivity Δ as $M(D) = \theta(D) + \text{Noise}$, where Noise is sampled from a Gaussian distribution $\mathcal{N}(0, \Delta^2/\mu^2)$. Then, M is μ -GDP.*

If M_1, M_2, \dots, M_k are k GDP mechanisms with parameters $\mu_1, \mu_2, \dots, \mu_k$, then the following GDP composition theorem holds:

Theorem 4. *The k -fold composition of μ_i -GDP mechanisms is $\sqrt{\mu_1^2 + \mu_2^2 + \dots + \mu_k^2}$ -GDP.*

Finally, one can convert GDP guarantees to the standard (ϵ, δ) -DP guarantee using the following theorem:

Theorem 5. *A mechanism is μ -GDP if and only if it is (ϵ, δ) -DP where*

$$\delta = \phi\left(-\frac{\epsilon}{\mu} + \frac{\mu}{2}\right) - e^\epsilon \phi\left(-\frac{\epsilon}{\mu} - \frac{\mu}{2}\right)$$

Besides the mathematical elegance, a key advantage of GDP is that it provides a tighter analysis of composition guarantees of differentially private mechanisms.

2.3. Notation

For rest of the paper, we adopt the following notation. We denote by H_k the histogram for feature k and \hat{H}_k for the corresponding differentially private histogram. We use K to denote the total number of features. By a slight abuse of notation, we write $x_i \in H_k(b)$ to mean that the data point x_i belongs to the histogram bin b , and use $\rho(H_k, x_i)$ to look up the bin b such that $x_i \in H_k(b)$.

3. Algorithms

To add differential privacy guarantees to EBMs, we modify the EBM training procedure Algorithm 1, yielding DP-EBM Algorithm 2. We first modify the pre-processing procedure to generate differentially private bins (published as histograms \hat{H}_k per feature), which log the bin ranges and how many data points fall in each bin (line 7). Next, we analyze the boosting process. In traditional tree building, there are two major data-intensive operations: learning the structure of the tree (what feature and feature threshold to install at each node in the tree), and calculating the predicted value of each leaf node (Breiman et al., 1984). Prior work on differentially private tree learning typically splits budget between choosing which features to split on, where to split

them, and learning prediction values for each leaf node (Fletcher & Islam, 2015; Wang et al.; Li et al., 2020).

EBMs naturally avoid spending any privacy budget on choosing which features to include in each tree – the “round-robin” schedule of visiting features is completely data agnostic. Furthermore, by choosing the splitting thresholds at random, we can learn the entire structure of each tree without looking at any training data (line 14). Prior work and our empirical evaluations both show that choosing random splits results in little accuracy loss (Geurts et al., 2006; Fletcher & Islam, 2019). We therefore spend the entirety of our budget per iteration on learning the values for each leaf node, which are simply averages of the residuals for the data belonging to each node (lines 16-18). For each leaf, we sum the residuals of data belonging to that leaf, add calibrated Gaussian noise based on their bounded sensitivity R , and divide by a differentially private count of data in each leaf (contained in the previously published \hat{H}_k). Then, as in non-private EBMs, the noisy tree is merged into the feature function f_k (line 20), and the cyclic boosting procedure moves onto the next feature and continues for all E epochs. The pseudocode for the DP-EBM algorithm is described in Algorithm 2, with modifications to the non-private version highlighted in blue.

We now provide the privacy analysis of our algorithm using the GDP framework. Our proof of privacy has the following two components. First we fix a single iteration of the algorithm (lines 13-26), and show that each iteration is $\frac{1}{\sigma} - \text{GDP}$. At the end of each iteration, we publicly release the functions f_k^t for all k . Note that although the final model only uses the f_k^t values of the *last iteration*, releasing every f_k^t leads to a simpler privacy analysis. Next, we calculate the total privacy loss of our algorithm by simply viewing it as a composition of $E \cdot K$ private mechanisms. It is important to note that composition theorems work even when the mechanisms depend on each other.

Theorem 6. *Each iteration of our algorithm is $\frac{1}{\sigma} - \text{GDP}$.*

Proof. We observe that calculating T in the line 16 is the only step of our algorithm where we access the sensitive information of the users. Thus to prove the theorem we need to argue that the noise we are adding satisfies requirements of Theorem 3. Consider

$$\begin{aligned} T &= \eta \cdot \sum_{b \in S_\ell} \sum_{x_i \in \hat{H}_k(b)} r_i^t \\ &= \eta \cdot \sum_{b \in S_\ell} \sum_{x_i \in \hat{H}_k(b)} \left(y_i - \sum_k f_k^{t-1}(\rho(\hat{H}_k, x_i)) \right) \\ &= \eta \cdot \left(\sum_{b \in S_\ell} \sum_{x_i \in \hat{H}_k(b)} y_i \right) - Z \end{aligned}$$

The second equality follows from the definition of r_i^t as given in the line 25 of the algorithm. Further, Z is computed using publicly released J_k^{t-1} values from the iteration $t - 1$, and hence does not depend on the user data. Therefore, the amount of noise we need to add to the statistic T depends on the sensitivity of quantity $\left(\sum_{b \in S_\ell} \sum_{x_i \in \hat{H}_k(b)} y_i\right)$, which we argue is bounded by at most R . This follows from three simple facts: 1) The range of each y_i is bounded by at most R ; 2) For each feature, each user contributes exactly to one bin of the histogram; 3) Random splits performed in line 14 of our algorithm partition the histogram bins into disjoint splits, hence each users’ data belongs to precisely one split. The proof now follows from Theorem 3. \square

Theorem 7. *Our algorithm is $\frac{\sqrt{E \cdot K}}{\sigma} - GDP$.*

Proof. As each iteration of our algorithm is $\frac{1}{\sigma} - GDP$, the proof follows from the composition of GDP-mechanisms as given in Theorem 4 over all $E \cdot K$ iterations. \square

The GDP bounds can be converted into (ϵ, δ) -DP guarantees using Theorem 5. To calibrate σ in line 17, we fix the ϵ and δ privacy parameters we want to achieve, use Theorem 5 to calculate μ , and finally calculate σ by setting $\mu = \frac{\sqrt{E \cdot K}}{\sigma}$.

4. Experiments

We compare the following algorithms on four classification and four regression datasets:

- DP-EBM (Algorithm 2): We use the following (default) parameters for all experiments: `max_bins = 32`, `learning_rate = 0.01`, `n_epochs = 300`, `max_leaves = 3`, with 10% of the total privacy budget allocated to binning and 90% to training. We present two results for DP-EBMs: “EBM-Classic”, where we apply strong composition from (Kairouz et al., 2017), and “EBM-GDP”, where composition is more tightly tracked via Gaussian Differential Privacy (Dong et al., 2019).
- Generalized Linear Models: Linear and Logistic Regression are widely used methods for interpretable machine learning. For both models, we use IBM’s differential privacy library (Holohan, 2019) which follows the algorithms described in (Sheffet, 2015; Imtiaz & Sarwate, 2016) for linear regression and in (Chaudhuri et al., 2011) for logistic regression.
- DP Boost: DPBoost is a differentially private gradient boosted decision tree algorithm introduced by (Li et al., 2020). DPBoost builds on top of LightGBM, a popular tree-based boosting framework (Ke et al., 2017).

To evaluate performance, we generate 25 randomly drawn 80/20 train-test splits and report the average test-set accuracy and standard deviation at varying ϵ and fixed $\delta = 10^{-6}$. Results are presented in Tables 2 and 3 using root mean squared error (RMSE) as the metric for regression and area under the ROC curve (AUROC) for classification.

All models were trained using default or recommended parameters from the literature or open source repositories. Hyperparameter tuning is a privacy-sensitive operation, and how to best partition budget and tune parameters of differentially private models is an open research problem (Liu & Talwar, 2019; Kusner et al., 2015). We did not tune hyperparameters to avoid this complexity. The default parameters for DP-EBMs appear to work well on a variety of datasets, which helps conserve the privacy budget and also makes DP-EBMs easy to use in practice.

The datasets used in these experiments (with the exception of the healthcare data, which contains real patient data) are publicly available and summarized in Table 1. We include results from the private healthcare dataset to highlight how these models might perform in a high stakes setting where both privacy and interpretability are critical.

Table 1. Experiment dataset statistics and descriptions.

Dataset	Domain	N	K	Task	% Pos
Adult Income	Finance	32,561	14	Clas	24.1%
Telco Churn	Business	7,043	19	Clas	26.6%
Credit	Finance	284,807	30	Clas	0.2%
Healthcare	Medicine	14,199	46	Clas	10.9%
Cal-Housing	Real Estate	20,499	8	Reg	–
Elevators	Systems	16,598	18	Reg	–
Pol	Business	15,000	49	Reg	–
Wine-Quality	Food	5,297	11	Reg	–

4.1. Discussion of Experimental Results

As shown in Figure 1 and again in more detail in Tables 2 and 3, DP-EBMs exhibit strong performance across a wide range of ϵ values. In classification, the difference in AUROC between non-private DP-EBMs and even $\epsilon = 0.5$ DP-EBMs is at most 0.05, which is a comparatively modest tradeoff for strong privacy guarantees. At a more modest $\epsilon = 4$, the average AUROC for DPBoost, Logistic Regression and DP-EBMs is 0.62, 0.56 and 0.88, respectively. The datasets chosen are not particularly favorable for EBMs – in the non-private setting, LightGBM outperforms EBMs in over half of our experiments. However, when training with differential privacy, DP-EBMs outperform other models in all 40 experimental settings. In the following sections, we offer some thoughts on why the DP-EBM algorithm might be comparatively well suited for differential privacy.




Figure 3. Risk of dying as a function of age from three EBMs trained on the healthcare dataset with varying privacy guarantees.

4.1.1. DP-EBMS VS LINEAR MODELS

In the case of DP linear and logistic regression – the current standard for intelligible and private learning – we believe the additional accuracy of DP-EBMs might be explained by the extra flexibility of the GAM model class. In the non-private setting, the non-linear functions learned by GAMs often result in a boost in accuracy over the linear functions learned by linear models, and this advantage appears to translate to the private setting as well. In addition, the iterative nature of gradient boosting might give DP-EBMs the ability to recover from the impact of noise earlier in training.

4.1.2. DP-EBMS VS DPBOOST

Although it may not be surprising that DP-EBMs outperform restricted models such as DP linear regression, it is a little surprising that DP-EBMs outperform DPBoost which is a less restricted model class than DP-EBMs. We believe this might be due to the significant privacy budget savings when learning each tree. Unlike other DP tree-based learning algorithms, DP-EBMs spend no budget learning the structure of each tree, and focus exclusively on learning the best leaf node values. In addition, by growing shallow trees, each leaf often contains a large proportion of the dataset – with default parameters of 2 random splits, in expectation each leaf contains $\approx \frac{1}{3}$ of the full data. This ensures that the impact of noise in the differentially private average calculated per iteration is dispersed across many training datapoints. In contrast, each tree in LightGBM/DPBoost contains up to 31 leaf nodes, thereby operating on much smaller amounts of data and magnifying the impact of noise.

4.1.3. DP-EBM: CLASSIC COMPOSITION VS GDP

We also compare two variants of privacy budget tracking in DP-EBMs: “EBM-Classic”, which uses strong composition from (Kairouz et al., 2017), and “EBM-GDP” which uses

Gaussian differential privacy recently proposed by (Dong et al., 2019). While we can show analytically that budget tracking with GDP is tighter for our algorithm (and therefore requires less noise for the same privacy loss), it is interesting that the differences in final model accuracy are typically only noticeable in strong privacy settings ($\epsilon \leq 1$).

5. Discussion

5.1. Editing DP-EBM Models

While this paper has primarily focused on introducing and comparing DP-EBMs in terms of standard machine learning metrics, we believe it is important to highlight the unique capabilities that arise when combining interpretability with differential privacy. For example, recent work has shown that adding differentially private noise to machine learning algorithms can disproportionately impact minority groups (Cummings et al., 2019; Bagdasaryan & Shmatikov, 2019). This concern is compounded when models are trained and deployed in high-risk domains like finance and healthcare – even small errors on sparse regions of the feature space can have disastrous consequences.

With intelligible models like DP-EBMs, some effects of noise on predictions are visible. Figure 3 shows the shape function DP-EBMs learned for the “Age” feature in the healthcare dataset at 3 different levels of privacy. In many healthcare problems, risk should monotonically increase as a function of age. While the non-private and $\epsilon = 4$ models learn this behavior, there is significant distortion as a result of differentially private noise at $\epsilon = 1$. In this example, the $\epsilon = 1$ model suggests that patients who are 80 are considerably lower risk than those who are 77 or 82, which is almost certainly not a real signal in the dataset. By using an intelligible model, domain experts can inspect the shape function f_{age} and prevent deploying a risky model that under predicts on 80 year olds.

Differentially Private EBMs

Table 2. Area Under the ROC Curve (AUROC) algorithm comparison on classification datasets. Higher is better.

DATASET	ϵ	DPBOOST	LOGISTIC REGRESSION	DPEBM-CLASSIC	DPEBM-GDP
ADULT-INCOME	0.5	0.558 \pm 0.045	0.488 \pm 0.125	0.873 \pm 0.007	0.875 \pm 0.005
	1.0	0.566 \pm 0.034	0.471 \pm 0.111	0.880 \pm 0.006	0.883 \pm 0.005
	2.0	0.629 \pm 0.045	0.521 \pm 0.109	0.886 \pm 0.005	0.887 \pm 0.004
	4.0	0.734 \pm 0.019	0.549 \pm 0.068	0.889 \pm 0.004	0.889 \pm 0.004
	8.0	0.805 \pm 0.011	0.534 \pm 0.070	0.890 \pm 0.004	0.890 \pm 0.004
	<i>Non-Private</i>	0.928 \pm 0.003	0.603 \pm 0.066	0.923 \pm 0.003	0.923 \pm 0.003
CREDIT-FRAUD	0.5	0.442 \pm 0.138	0.558 \pm 0.076	0.959 \pm 0.015	0.966 \pm 0.012
	1.0	0.438 \pm 0.114	0.544 \pm 0.135	0.965 \pm 0.014	0.966 \pm 0.013
	2.0	0.467 \pm 0.101	0.526 \pm 0.118	0.969 \pm 0.011	0.969 \pm 0.011
	4.0	0.465 \pm 0.142	0.539 \pm 0.172	0.969 \pm 0.011	0.969 \pm 0.011
	8.0	0.556 \pm 0.145	0.546 \pm 0.156	0.969 \pm 0.011	0.969 \pm 0.011
	<i>Non-Private</i>	0.726 \pm 0.099	0.922 \pm 0.019	0.965 \pm 0.011	0.965 \pm 0.011
HEALTHCARE	0.5	0.515 \pm 0.054	0.463 \pm 0.081	0.714 \pm 0.036	0.793 \pm 0.018
	1.0	0.505 \pm 0.051	0.479 \pm 0.071	0.789 \pm 0.016	0.818 \pm 0.012
	2.0	0.499 \pm 0.046	0.495 \pm 0.081	0.822 \pm 0.012	0.830 \pm 0.011
	4.0	0.567 \pm 0.047	0.542 \pm 0.038	0.834 \pm 0.011	0.835 \pm 0.010
	8.0	0.638 \pm 0.036	0.529 \pm 0.048	0.836 \pm 0.010	0.837 \pm 0.010
	<i>Non-Private</i>	0.836 \pm 0.011	0.744 \pm 0.014	0.847 \pm 0.010	0.847 \pm 0.010
TELCO-CHURN	0.5	0.484 \pm 0.100	0.541 \pm 0.227	0.812 \pm 0.020	0.829 \pm 0.014
	1.0	0.458 \pm 0.088	0.479 \pm 0.239	0.832 \pm 0.013	0.835 \pm 0.011
	2.0	0.534 \pm 0.109	0.527 \pm 0.236	0.837 \pm 0.010	0.838 \pm 0.012
	4.0	0.716 \pm 0.067	0.615 \pm 0.138	0.838 \pm 0.011	0.839 \pm 0.011
	8.0	0.787 \pm 0.014	0.673 \pm 0.105	0.839 \pm 0.011	0.839 \pm 0.011
	<i>Non-Private</i>	0.836 \pm 0.008	0.844 \pm 0.010	0.848 \pm 0.009	0.848 \pm 0.009

Table 3. Root Mean Squared Error (RMSE) algorithm comparison on regression datasets. Lower is better.

DATASET	ϵ	DPBOOST	LINEAR REGRESSION	DPEBM-CLASSIC	DPEBM-GDP
CAL-HOUSING	0.5	383072 \pm 41952	111967 \pm 1080	85652 \pm 2724	79967 \pm 1929
	1.0	204277 \pm 19350	110241 \pm 1101	78527 \pm 1230	76827 \pm 1470
	2.0	122494 \pm 7066	109518 \pm 1244	75491 \pm 1404	74573 \pm 1152
	4.0	96336 \pm 3043	108882 \pm 1370	73967 \pm 1028	73754 \pm 1022
	8.0	90029 \pm 2508	107815 \pm 1460	73327 \pm 1118	73165 \pm 955
	<i>Non-Private</i>	47007 \pm 885	69850 \pm 1164	51644 \pm 925	51644 \pm 925
ELEVATORS	0.5	0.051 \pm 0.005	4.671 \pm 1.975	0.010 \pm 0.001	0.006 \pm 0.000
	1.0	0.025 \pm 0.002	2.669 \pm 1.214	0.007 \pm 0.000	0.005 \pm 0.000
	2.0	0.013 \pm 0.001	1.384 \pm 0.570	0.006 \pm 0.000	0.005 \pm 0.000
	4.0	0.008 \pm 0.000	0.754 \pm 0.202	0.005 \pm 0.000	0.004 \pm 0.000
	8.0	0.006 \pm 0.000	0.410 \pm 0.201	0.004 \pm 0.000	0.004 \pm 0.000
	<i>Non-Private</i>	0.002 \pm 0.000	0.003 \pm 0.000	0.004 \pm 0.000	0.004 \pm 0.000
POL	0.5	78.190 \pm 9.583	31.326 \pm 0.418	35.156 \pm 1.728	30.988 \pm 0.962
	1.0	50.527 \pm 5.482	30.640 \pm 0.288	30.911 \pm 1.014	28.391 \pm 0.585
	2.0	47.511 \pm 4.636	30.500 \pm 0.248	27.616 \pm 0.644	26.303 \pm 0.561
	4.0	45.592 \pm 2.942	30.463 \pm 0.256	25.454 \pm 0.389	24.934 \pm 0.332
	8.0	45.435 \pm 1.109	30.459 \pm 0.258	24.625 \pm 0.230	24.313 \pm 0.237
	<i>Non-Private</i>	4.703 \pm 0.228	30.464 \pm 0.264	13.780 \pm 0.667	13.780 \pm 0.667
WINE-QUALITY	0.5	4.647 \pm 0.390	3.621 \pm 1.740	1.589 \pm 0.132	0.938 \pm 0.036
	1.0	2.151 \pm 0.302	2.133 \pm 0.713	1.181 \pm 0.074	0.841 \pm 0.025
	2.0	1.299 \pm 0.092	1.263 \pm 0.322	0.935 \pm 0.042	0.779 \pm 0.018
	4.0	0.946 \pm 0.043	0.940 \pm 0.100	0.807 \pm 0.019	0.746 \pm 0.011
	8.0	0.847 \pm 0.021	0.839 \pm 0.035	0.751 \pm 0.013	0.733 \pm 0.014
	<i>Non-Private</i>	0.622 \pm 0.013	0.759 \pm 0.015	0.681 \pm 0.012	0.681 \pm 0.012




Figure 4. (Left) Original learned f_{age} from the healthcare dataset with $\epsilon = 1$. (Middle, Right) Postprocessed to enforce monotonicity.

In addition to catching errors, domain experts can also correct unwanted learned effects without impacting privacy. Because the shape function f_{age} exactly describes how a model makes predictions, users can edit graphs of any feature to change the model. In our example, modifying y-axis value for f_{age} at ages 78-81 to remove the unwanted blip would remove this noise bias from the model. This form of model editing uses no data, and therefore results in no additional privacy loss under the post-processing property of differential privacy (Dwork et al., 2006). The ability to safely inspect and edit DP-EBM models before deployment is important for creating trust in differentially private models in high risk situations because some of the impacts of noise can be corrected.

5.2. Constraints such as monotonicity

More complex forms of editing are also possible. For example, we can ensure monotonicity across the entire feature by borrowing from the calibration literature and applying *isotonic regression* on the shape function (Chakravarti, 1989). Isotonic regression uses the Pool Adjacent Violators (PAV) algorithm to minimize the edits necessary to ensure monotonicity, and is optimal with respect to squared error on the differences between the two functions (vanEeden, 1958). Importantly, this process only uses public information from DP-EBMs – the learned shape functions f_k , and the histogram definition \hat{H}_k .

Figure 4 shows the effects of applying isotonic regression to the noisy f_{age} . While enforcing monotonicity is possible with models more expressive than DP-EBMs, this typically requires additional constraints during training and may consume more privacy budget and complicate the privacy analysis. It is a nice advantage of DP-EBMs that monotonicity can be achieved cleanly via post-processing.

5.3. Differential Privacy as a Regularizer

Figures 3 and 5 also show that adding modest amounts of differentially private noise, like $\epsilon = 4$, can act as a regularizer to the model. The rise in risk between ages 50 and 90 here is smoother, whereas the non-private version learns a “jumpy” function. Smoothness is traditionally a difficult property to achieve with tree-based boosted GAMs. The non-private EBM algorithm typically wraps the training process in multiple iterations of bagging to make graphs smoother (Lou et al., 2012; Caruana et al., 2015). Our experiments suggest that modest amounts of differentially private noise might act as an effective regularization tool.

The relationship between smoothness and interpretability is complex: smooth graphs may be easier to interpret, but over-regularization can hide real signals in the data. The use of differential privacy as a regularizer is well known (Chaudhuri et al., 2011). Our paper visibly reinforces this notion through intelligibility, and we believe studying this effect further on GAMs might be interesting future research.

6. Conclusion

We present DP-EBMs, a differentially private learning algorithm for GAMs which achieves remarkably high accuracy and interpretability with strong privacy guarantees. Our empirical evaluations show that DP-EBMs outperform other differentially private learning algorithms for both classification and regression on tabular datasets. Beyond just accuracy, we also show how interpretability can complement differential privacy by enabling users to uncover undesirable effects of noise, edit unwanted bias out of their models, and enforce desirable constraints like monotonicity with no additional privacy loss. These practical advantages might represent an important step forward for enabling the use of differentially private models in industries like healthcare, finance, and criminal justice.

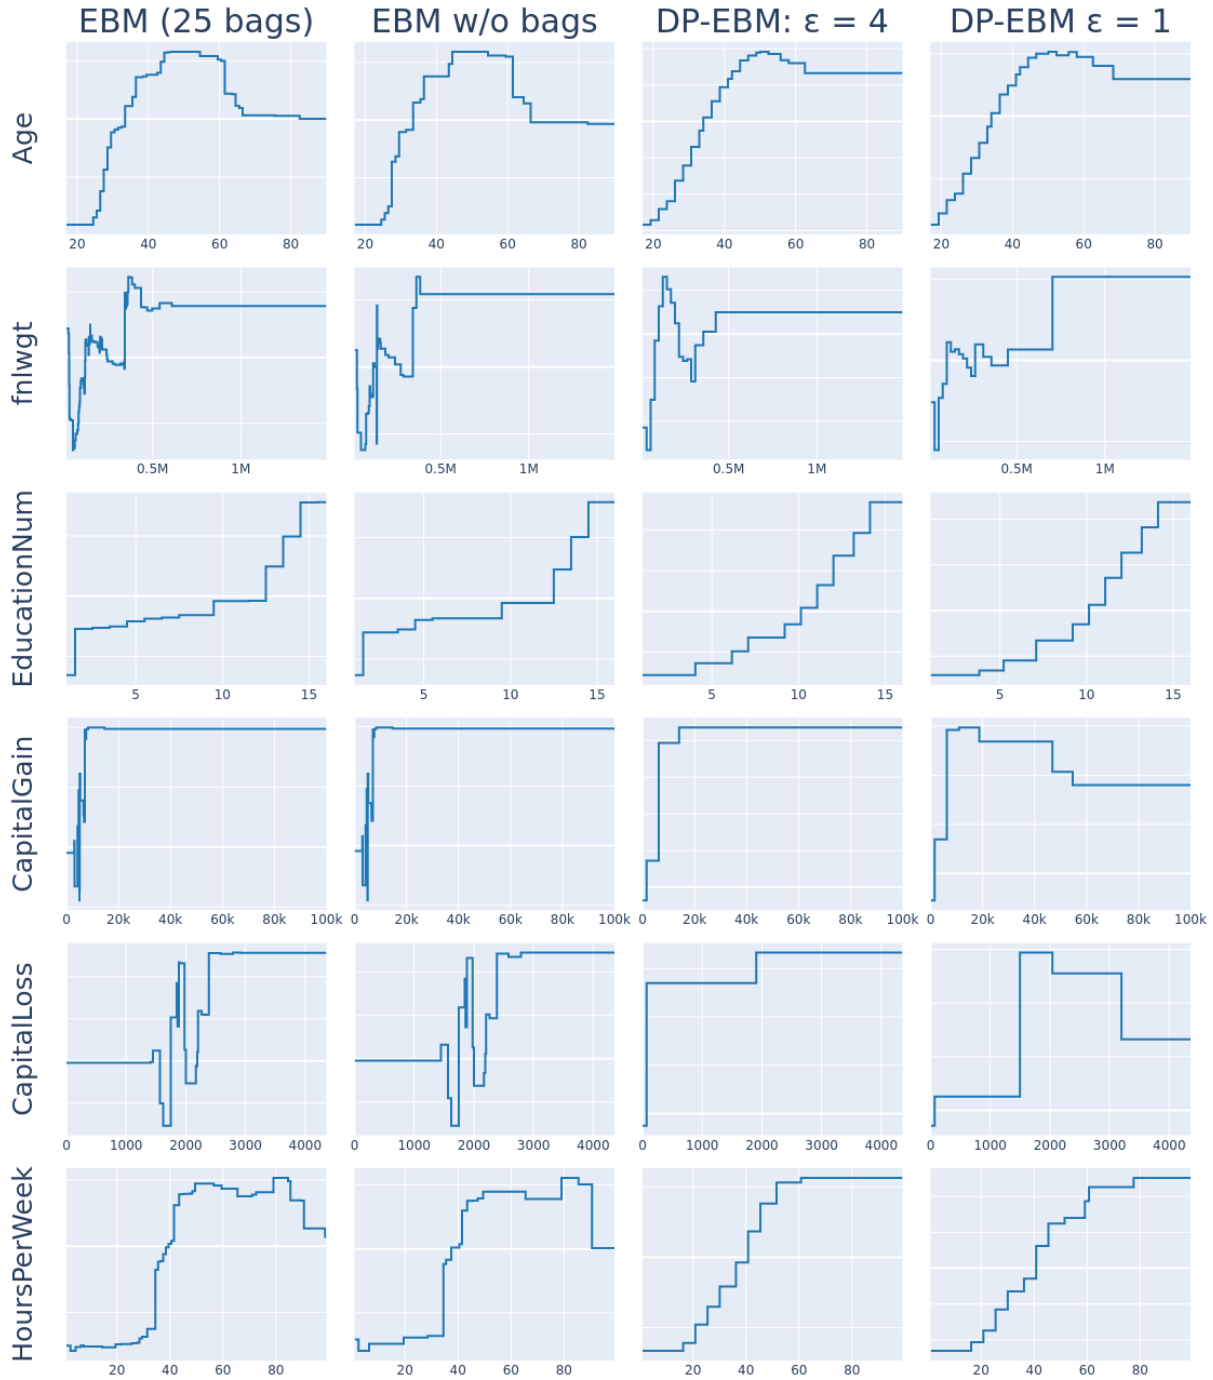


Figure 5. Shape function comparisons for all numeric features in the Adult Income dataset. We include the standard EBMs wrapped in 25 layers of bagging, EBMs without bagging, and two DP-EBMs at different privacy levels ($\epsilon = 4$ and $\epsilon = 1$). As expected, adding Gaussian DP noise acts as a strong regularizer — the graphs on the right are smoother than those on the left. In some cases this regularization is too strong, yet in other cases such as *fnlwgt* the extra regularization might actually reduce model variance and improve intelligibility.

References

- Bagdasaryan, E. and Shmatikov, V. Differential privacy has disparate impact on model accuracy. *arXiv preprint arXiv:1905.12101*, 2019.
- Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. *Classification and regression trees*. CRC press, 1984.
- Bu, Z., Dong, J., Long, Q., and Su, W. J. Deep learning with gaussian differential privacy. *Harvard data science review*, 2020(23), 2020.
- Buitinck, L., Louppe, G., Blondel, M., Pedregosa, F., Mueller, A., Grisel, O., Niculae, V., Prettenhofer, P., Gramfort, A., Grobler, J., Layton, R., VanderPlas, J., Joly, A., Holt, B., and Varoquaux, G. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, pp. 108–122, 2013.
- Carlini, N., Liu, C., Kos, J., Erlingsson, Ú., and Song, D. The secret sharer: Measuring unintended neural network memorization & extracting secrets. 2018.
- Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., and Elhadad, N. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1721–1730, 2015.
- Chakravarti, N. Isotonic median regression: a linear programming approach. *Mathematics of operations research*, 14(2):303–308, 1989.
- Chang, C.-H., Tan, S., Lengerich, B., Goldenberg, A., and Caruana, R. How interpretable and trustworthy are gams? *arXiv preprint arXiv:2006.06466*, 2020.
- Chaudhuri, K., Monteleoni, C., and Sarwate, A. D. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.
- Chen, T. and Guestrin, C. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.
- Cummings, R., Gupta, V., Kimpara, D., and Morgenstern, J. On the compatibility of privacy and fairness. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, pp. 309–315, 2019.
- Dong, J., Roth, A., and Su, W. J. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019.
- Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284. Springer, 2006.
- Dwork, C., Roth, A., et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Fletcher, S. and Islam, M. Z. A differentially private decision forest. *AusDM*, 15:99–108, 2015.
- Fletcher, S. and Islam, M. Z. Decision tree classification with differential privacy: A survey. *ACM Computing Surveys (CSUR)*, 52(4):1–33, 2019.
- Geurts, P., Ernst, D., and Wehenkel, L. Extremely randomized trees. *Machine learning*, 63(1):3–42, 2006.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- Hastie, T. J. and Tibshirani, R. J. *Generalized additive models*, volume 43. CRC press, 1990.
- Holohan, N. Diffprivlib: The ibm differential privacy library. <https://github.com/IBM/differential-privacy-library>, 2019.
- Imtiaz, H. and Sarwate, A. D. Symmetric matrix perturbation for differentially-private principal component analysis. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2339–2343. IEEE, 2016.
- Inc., P. T. Collaborative data science, 2015. URL <https://plot.ly>.
- Kairouz, P., Oh, S., and Viswanath, P. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in neural information processing systems*, pp. 3146–3154, 2017.

- Kusner, M., Gardner, J., Garnett, R., and Weinberger, K. Differentially private bayesian optimization. In *International conference on machine learning*, pp. 918–927. PMLR, 2015.
- Li, Q., Wu, Z., Wen, Z., and He, B. Privacy-preserving gradient boosting decision trees. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 784–791, 2020.
- Liu, J. and Talwar, K. Private selection from private candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 298–309, 2019.
- Lou, Y., Caruana, R., and Gehrke, J. Intelligible models for classification and regression. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 150–158, 2012.
- Lundberg, S. M. and Lee, S.-I. A unified approach to interpreting model predictions. In *Advances in neural information processing systems*, pp. 4765–4774, 2017.
- Melis, L., Song, C., De Cristofaro, E., and Shmatikov, V. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 691–706. IEEE, 2019.
- Nori, H., Jenkins, S., Koch, P., and Caruana, R. Interpretml: A unified framework for machine learning interpretability. *arXiv preprint arXiv:1909.09223*, 2019.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.
- Sheffet, O. Private approximations of the 2nd-moment matrix using existing techniques in linear regression. *arXiv preprint arXiv:1507.00056*, 2015.
- vanEeden, C. Testing and estimating ordered parameters of probability distribution. 1958.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D.,
- Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- Wang, C., Han, B., Patel, B., Mohideen, F., and Rudin, C. In pursuit of interpretable, fair and accurate machine learning for criminal recidivism prediction. *arXiv preprint arXiv:2005.04176*, 2020.
- Wang, K. W., Dick, T., and Balcan, M.-F. Scalable and provably accurate algorithms for differentially private distributed decision tree learning.