
Treeffuser: Probabilistic Predictions via Conditional Diffusions with Gradient-Boosted Trees

Nicolas Beltran-Velez*¹ Alessandro Antonio Grande*^{2,3} Achille Nazaret*^{1,3}
Alp Kucukelbir^{1,4} David Blei^{1,5}

¹Department of Computer Science, Columbia University, New York, USA

²Halvorsen Center for Computational Oncology, Memorial Sloan Kettering Cancer Center, New York, USA

³Irving Institute for Cancer Dynamics, Columbia University, New York, USA

⁴Fero Labs, New York, USA

⁵Department of Statistics, Columbia University, New York, USA

Abstract

Probabilistic prediction aims to compute predictive distributions rather than single point predictions. These distributions enable practitioners to quantify uncertainty, compute risk, and detect outliers. However, most probabilistic methods assume parametric responses, such as Gaussian or Poisson distributions. When these assumptions fail, such models lead to bad predictions and poorly calibrated uncertainty. In this paper, we propose Treeffuser, an easy-to-use method for probabilistic prediction on tabular data. The idea is to learn a conditional diffusion model where the score function is estimated using gradient-boosted trees. The conditional diffusion model makes Treeffuser flexible and non-parametric, while the gradient-boosted trees make it robust and easy to train on CPUs. Treeffuser learns well-calibrated predictive distributions and can handle a wide range of regression tasks—including those with multivariate, multimodal, and skewed responses. We study Treeffuser on synthetic and real data and show that it outperforms existing methods, providing better calibrated probabilistic predictions. We further demonstrate its versatility with an application to inventory allocation under uncertainty using sales data from Walmart. We implement Treeffuser in <https://github.com/blei-lab/treeffuser>.

1 Introduction

In this paper, we develop a new method for probabilistic prediction from tabular data. This problem is important to many fields, such as power generation [1], finance [2], and healthcare [3]. It drives decision processes such as supply chain planning [4], risk assessment [5, 6], and policy-making [7].

Example: Manufacturing plants measure information as they operate. This information—properties of raw materials, operational flows, temperatures, and level measurements—collectively determine the output of the plant [8]. From this data, how can manufacturers adapt operations to reduce emissions while maximizing profits? The answer requires both good predictions and estimates of uncertainty, so as to trade off the risk of failure with the reward of lower emissions and higher profit. More broadly, such industrial workflow problems often rely on predictions from vast amounts of tabular data—observations of variables arranged in a table [9, 10].

*Equal contribution, authors listed in alphabetical order.

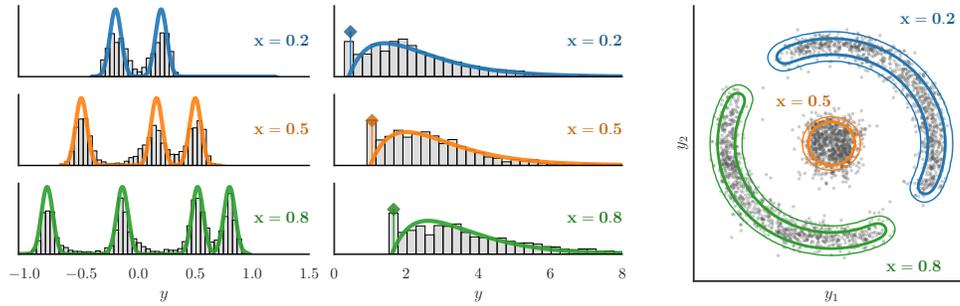


Figure 1: Samples $\mathbf{y} \mid \mathbf{x}$ from Treeffuser vs. true densities, for multiple values of \mathbf{x} under three different scenarios. Treeffuser captures arbitrarily complex conditional distributions that vary with \mathbf{x} .

Our method builds on two ideas: diffusions [11] and gradient-boosted trees [12]. Diffusions accurately estimate conditional distributions, but existing methods are not designed for tabular data. Decision trees excel at analyzing tabular data, but do not provide probabilistic predictions. Our method, which we call *Treeffuser*, combines the advantages of both formalisms. It defines an accurate tree-based diffusion model for probabilistic prediction which can easily handle large datasets.

Fig. 1 shows Treeffuser in action. We feed it data from three complex distributions—a multimodal response with varying components, an inflated response with shifting support, and a multivariate response with dynamic correlations. For each task, Treeffuser uses trees to learn a diffusion model and then outputs samples from the target conditional. These samples fully capture the complexity of the response functions.

Treeffuser exhibits several advantages:

- It is nonparametric. It makes few assumptions about the form of the conditional response distribution. For example, it can estimate response distributions that are multivariate, multimodal, skewed, and heavy-tailed.
- It is efficient. Diffusions can be slow to train [13]. Treeffuser relies on trees and is fast. For instance, in section 5.3, Treeffuser trains from a table with 112,000 observations and 27 variables in 53 seconds on a laptop.
- It is accurate. On benchmark datasets, Treeffuser outperforms NGBoost [14], IBUG [15], quantile regression [16] and deep ensembles [17]. It provides better probabilistic predictions, including more precise quantile estimations and accurate mean predictions.

The rest of the paper is organized as follows. Section 2 reviews diffusions and gradient-boosted trees. Section 3 integrates these two ideas into Treeffuser and justifies the method. Section 4 discusses related work. Section 5 studies synthetic and standard benchmark datasets.

2 Background

Treeffuser combines two ideas: diffusion models and gradient-boosted trees (GBTs). This section provides a gentle introduction to both topics. Familiar readers can skip ahead to Section 3.

2.1 Diffusion Models

A diffusion model is a generative model that learns an arbitrarily complex distribution $\pi(\mathbf{y})$. It consists of two processes: forward and reverse diffusion.

The forward process takes the target distribution $\pi(\mathbf{y})$ and continuously transforms it into a simple distribution. It does so by progressively adding noise to samples from π :

$$\begin{cases} \text{Draw } \mathbf{y}(0) \sim \pi, \\ \text{Evolve } \mathbf{y}(t) \text{ until } T \text{ according to: } d\mathbf{y}(t) = f(\mathbf{y}(t), t)dt + g(t)d\mathbf{w}(t), \end{cases} \quad (1)$$

where Eq. (1) defines a stochastic differential equation (SDE) with a standard Brownian motion $\mathbf{w}(t)$, and drift and diffusion functions f and g . The time horizon T is large, such that the resulting marginal distribution $p_{\text{simple}} := p_T(\mathbf{y}(T)) = \int_{\mathbf{y}'} p_T(\mathbf{y}(T) | \mathbf{y}(0) = \mathbf{y}') \pi(\mathbf{y}') d\mathbf{y}'$ is agnostic to π .

The reverse process transforms the simple distribution back into the target distribution by denoising perturbed samples from p_{simple} . It posits the following model, which runs backward from T to 0:

$$\begin{cases} \tilde{\mathbf{y}}(T) \sim p_{\text{simple}}, \\ d\tilde{\mathbf{y}}_t = [f(\tilde{\mathbf{y}}(t), t) - g(t)^2 \nabla_{\tilde{\mathbf{y}}(t)} \log p_t(\tilde{\mathbf{y}}(t))] dt + g(t) d\tilde{\mathbf{w}}(t), \end{cases} \quad (2)$$

where $\tilde{\mathbf{w}}(t)$ is a standard Brownian with reversed time. Anderson [18] shows that $\tilde{\mathbf{y}}(t) \stackrel{d}{=} \mathbf{y}(t)$ for each time t , and so $\tilde{\mathbf{y}}(0) \sim \pi$. This means that by drawing a sample from p_{simple} and then numerically approximating the reverse process in Eq. (2), we obtain samples from π .

However, the score function, $\mathbf{y} \mapsto \nabla_{\mathbf{y}} \log p_t(\mathbf{y})$, is usually unknown. Vincent [19] shows it can be estimated from the trajectories of the forward process as the minimizer of the following objective:

$$\left[\mathbf{y} \mapsto \nabla_{\mathbf{y}} \log p_t(\mathbf{y}) \right] = \arg \min_{s \in \mathcal{S}} \mathbb{E}_t \mathbb{E}_{\pi} \mathbb{E}_{p_{0t}} \left[\left\| \nabla_{\mathbf{y}(t)} \log p_{0t}(\mathbf{y}(t) | \mathbf{y}(0)) - s(\mathbf{y}(t), t) \right\|^2 \right], \quad (3)$$

where \mathcal{S} is the set of all possible functions of \mathbf{y} indexed by time t , and p_{0t} denotes the conditional distribution of $\mathbf{y}(t)$ given $\mathbf{y}(0)$. In practice, we set $t \sim \text{Uniform}([0, 1])$ and choose the drift f and diffusion function g so that p_{0t} is Gaussian. We detail our choice of f and g in Appendix C.

We approximate the expectations in Eq. (3) with the empirical distribution of the observed \mathbf{y} for \mathbb{E}_{π} , and with the known Gaussian trajectories $\mathbf{y}(t) | \mathbf{y}(0) = \mathbf{y}$ for $\mathbb{E}_{p_{0t}}$. The objective can then be minimized by parametrizing the score s with any function approximator, such as a neural network or, as we do, with trees. By estimating the score function, we effectively learn the distribution π .

For a more detailed introduction to diffusion models, we provide an expanded version of this section in Appendix A.

2.2 Gradient Boosted Trees

Consider the following common machine learning task. Given variables $(\mathbf{a}, b) \sim \pi$, where b is scalar, learn the function $F^*(\mathbf{a})$ that best approximates b from \mathbf{a} :

$$F^* = \arg \min_F \mathbb{E}_{(\mathbf{a}, b) \sim \pi} [L(F(\mathbf{a}), b)], \quad (4)$$

where L is a loss function measuring the quality of F , such as the squared loss $L(u, v) = (u - v)^2$. Gradient-boosted trees (GBTs) are a widely used non-parametric machine learning model for this task [12]. GBTs use decision trees [20] as simple building block functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ to form a good approximation \hat{F} of F^* . GBTs sequentially build approximations \hat{F}_i defined as

$$\hat{F}_i(\mathbf{a}) = \sum_{m=0}^{i-1} \varepsilon f_m(\mathbf{a}), \quad (5)$$

where $\varepsilon \in (0, 1)$ is a parameter akin to a learning rate.

Each decision tree f_i is constructed to minimize the squared error between the current approximation \hat{F}_i and the negative gradient of the loss function:

$$f_i = \arg \min_{f \in \text{Trees}} \mathbb{E}_{\mathbf{a}, b} \left[\left(f(\mathbf{a}) + \frac{\partial L(\hat{F}_i(\mathbf{a}), b)}{\partial \hat{F}_i(\mathbf{a})} \right)^2 \right]. \quad (6)$$

As the number of trees i increases, the approximation \hat{F}_i becomes a better minimizer of Eq. (4). Various modifications to this basic algorithm have been proposed, including different loss functions in Eq. (4), higher order optimizations for Eq. (6) and general heuristics for faster training [21–23].

3 Probabilistic Prediction via Treffuser

Treffuser tackles probabilistic prediction by learning a diffusion model with gradient-boosted trees. It is particularly well-suited to modeling tabular data. Trees offer useful inductive biases, natural

handling of categorical and missing data, and fast and robust training procedures. Diffusions eliminate the need for restrictive parametric families of distributions and protect against model misspecification.

Denote an independently distributed set of observations as $\mathcal{D} = \{(x_i, \mathbf{y}_i)\}_{i=1}^n$. Treeffuser forms the predictive distribution $\pi(\mathbf{y} \mid \mathbf{x})$ as a function of inputs \mathbf{x} . We first introduce conditional diffusion models and discuss the conditional score estimation problem for both univariate and multivariate outcomes. We then outline the training and sampling procedures for Treeffuser.

3.1 The Conditional Diffusion Model

Treeffuser produces a distribution over \mathbf{y} for each value of \mathbf{x} using conditional diffusion models. Unlike approaches that guide an unconditional model to achieve conditionality, Treeffuser follows the line of work that directly fits the conditional score function [24–26].

Conditional diffusion models. The diffusion models introduced in Section 2.1 target the marginal distribution of \mathbf{y} . Here, we extend them to conditional distributions $\pi_{\mathbf{x}}(\mathbf{y}) := \pi(\mathbf{y} \mid \mathbf{x})$.

To model $\pi_{\mathbf{x}}$, assign a diffusion process $\mathbf{y}_{\mathbf{x}}(t)$ to each value of \mathbf{x} . These processes share the same diffusion equation but have different boundary conditions corresponding to their target conditionals:

$$\begin{cases} \mathbf{y}_{\mathbf{x}}(0) \sim \pi_{\mathbf{x}}(\mathbf{y}), \\ d\mathbf{y}_{\mathbf{x}}(t) = f(\mathbf{y}_{\mathbf{x}}(t), t)dt + g(t)d\mathbf{w}(t). \end{cases} \quad (7)$$

As before, f and g are simple functions such that $\mathbf{y}_{\mathbf{x}}(T) \sim p_{\text{simple}}$ for all \mathbf{x} . Denote the marginal distribution of $\mathbf{y}_{\mathbf{x}}(t)$ as $p_{\mathbf{x},t}$. For two time points t and u , where $t > u$, denote the conditional distribution $\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(u)$ as p_{ut} . (Note how p_{ut} does not vary in \mathbf{x} .)

To match each \mathbf{x} -dependent forward SDE we have an \mathbf{x} -dependent reverse SDE of the form:

$$d\tilde{\mathbf{y}}_{\mathbf{x}}(t) = [f(\tilde{\mathbf{y}}_{\mathbf{x}}(t), t) - g(t)^2 \nabla_{\tilde{\mathbf{y}}_{\mathbf{x}}(t)} \log p_{\mathbf{x},t}(\tilde{\mathbf{y}}_{\mathbf{x}}(t))] dt + g(t)d\tilde{\mathbf{w}}(t), \quad (8)$$

where the score function $\nabla \log p_{\mathbf{x},t}$ now also depends on \mathbf{x} . Similar to unconditional diffusions, by estimating the *conditional* score, we can sample from $\mathbf{y}_{\mathbf{x}}$ by first sampling $\tilde{\mathbf{y}}_{\mathbf{x}}(T)$ from $p_{\text{simple}}(\mathbf{y})$ and then solving the corresponding reverse SDE.

Conditional score estimation objective. Estimating the conditional score follows a similar strategy as the unconditional version, but with the added requirement of simultaneously estimating the score for all \mathbf{x} . Recall that by Eq. (3), for a fixed \mathbf{x} , the function $s_{\mathbf{x}}^*$ defined by

$$s_{\mathbf{x}}^* = \arg \min_{s \in \mathcal{S}} \mathbb{E}_t \mathbb{E}_{\pi_{\mathbf{x}}(\mathbf{y}_{\mathbf{x}}(0))} \mathbb{E}_{p_{0t}(\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(0))} \left[\left\| \nabla_{\mathbf{y}_{\mathbf{x}}(t)} \log p_{0t}(\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(0)) - s(\mathbf{y}_{\mathbf{x}}(t), t) \right\|^2 \right]$$

satisfies $s_{\mathbf{x}}^*(\mathbf{y}_{\mathbf{x}}, t) = \nabla_{\mathbf{y}_{\mathbf{x}}(t)} \log p_{\mathbf{x},t}(\mathbf{y}_{\mathbf{x}}(t))$ for all \mathbf{y}, t and the fixed \mathbf{x} . Intuitively, if we allow $s \in \mathcal{S}$ to also take \mathbf{x} as input, we can gather all of these individual optimization problems into a single large problem by taking an additional expectation over $\mathbf{x} \sim \pi$, that is:

$$\arg \min_{S \in \mathcal{S}^+} \mathbb{E}_{\pi(\mathbf{x})} \mathbb{E}_t \mathbb{E}_{\pi(\mathbf{y}_{\mathbf{x}}(0))} \mathbb{E}_{p_{0t}} \left[\left\| \nabla_{\mathbf{y}_{\mathbf{x}}(t)} \log p_{0t}(\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(0)) - S(\mathbf{y}_{\mathbf{x}}(t), t, \mathbf{x}) \right\|^2 \right]. \quad (9)$$

Here, \mathcal{S}^+ represents the set of functions that take \mathbf{x} as an extra input along \mathbf{y} and t . The uppercase S emphasizes that S takes \mathbf{x} as input, contrary to lowercase s . The validity of this objective is given by the following result.

Theorem 1 (Optimal Conditional Objective). *Define S^* as the solution of Eq. (9). Then, for almost all $\mathbf{x}, \mathbf{y}, t$ with respect to $\pi(\mathbf{x}, \mathbf{y})$ and the Lebesgue measure on $t \in [0, T]$, we have*

$$S^*(\mathbf{y}, t, \mathbf{x}) = \nabla_{\mathbf{y}} \log p_{\mathbf{x},t}(\mathbf{y}). \quad (10)$$

We refer to Eq. (9) as the *conditional score objective*. The proof is provided in Appendix B.

3.2 The Trees

Treeffuser uses gradient-boosted trees (GBTs) to minimize Eq. (9). As GBTs work on scalar outputs, we separate the conditional score objective into an equivalent set of $d_{\mathbf{y}}$ independent scalar-valued

Algorithm 1: Treeffuser Training

Data: \mathcal{D} , R , SDE-specific (h, T)
Result: GBTs (u_1, \dots, u_d) optimizing Eq. (12)
 $\mathcal{D}' \leftarrow \emptyset$
for $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}) \in \mathcal{D}$ **do**
 for $r = 1, \dots, R$ **do**
 $t \sim \text{Uniform}[0, T]$
 $\zeta \sim \mathcal{N}(0, I_{d_y})$
 $\mathcal{D}' \leftarrow \mathcal{D}' \cup ((h(\zeta, t, \mathbf{y}^{(i)}), t, \mathbf{x}^{(i)}), -\zeta)$
for $k = 1, \dots, d_y$ **do**
 $\mathcal{D}^k \leftarrow \{(\mathbf{a}, \mathbf{b}_k) \mid (\mathbf{a}, \mathbf{b}) \in \mathcal{D}'\}$
 $U_k \leftarrow \text{TrainGBT}(\mathcal{D}^k)$
return (U_1, \dots, U_d)

Algorithm 2: Treeffuser Sampling

Data: GBTs (U_1, \dots, U_{d_y}) , input instance \mathbf{x} , discretization steps n_d , SDE-specific $(p_{\text{simple}}, T, f, g, \sigma)$
Result: A sample $\mathbf{y} \sim \pi(\mathbf{y} \mid \mathbf{x})$
 $\delta \leftarrow T/n_d$
 $\mathbf{y} \sim p_{\text{simple}}(\mathbf{y})$
 $t \leftarrow T$
for $i = 1, \dots, n_d$ **do**
 $\mathbf{w} \sim \mathcal{N}(0, I_{d_y})$
 $\tilde{f} \leftarrow f(\mathbf{y}, t) - g(t)^2 U(\mathbf{y}, t, \mathbf{x}) / \sigma(t; \mathbf{y})$
 $\mathbf{y} \leftarrow \mathbf{y} - (\tilde{f}\delta + g(t)\delta\mathbf{w})$
 $t \leftarrow t - \Delta$
return \mathbf{y}

sub-problems

$$S_k^* = \arg \min_{S_k \in \text{GBT}} \mathbb{E}_t \mathbb{E}_{\pi(\mathbf{x}, \mathbf{y})} \mathbb{E}_{p_{0t}(\mathbf{y}_x(t) \mid \mathbf{y})} \left[\left(\frac{\partial \log p_{0t}(\mathbf{y}_x(t) \mid \mathbf{y})}{\partial y_{\mathbf{x}}(t)_k} - S_k(\mathbf{y}_x(t), t, \mathbf{x}) \right)^2 \right], \quad (11)$$

where a_k denotes the k -th component of vector \mathbf{a} .

Recall that the drift and diffusion functions of the forward process are chosen such that p_{0t} is Gaussian. Let $m = m(t; \mathbf{y})$ and $\sigma = \sigma(t)$ denote the corresponding mean and standard deviation, respectively. Treeffuser replaces the partial derivative in Eq. (11) with its closed-form expression as a function of m and σ . Treeffuser further reparametrizes $S(\mathbf{y}, t, \mathbf{x})$ with $U(\mathbf{y}, t, \mathbf{x})/\sigma(t)$ and defines $h(\zeta, t, \mathbf{y}) = m(t; \mathbf{y}) + \zeta\sigma(t)$, the process by which a sample \mathbf{y} at time 0 gets diffused into a sample at time t with Gaussian noise ζ . The optimization problems in Eq. (11) are then:

$$\forall k \in \{1, \dots, d_y\}, \quad U_k^* = \arg \min_{U_k} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \pi} \mathbb{E}_t \mathbb{E}_{\zeta \sim \mathcal{N}(0, I_{d_y})} \left[(\zeta_k + U_k(h(\zeta, t, \mathbf{y}), t, \mathbf{x}))^2 \right]. \quad (12)$$

The next theorem justifies how the individual problems in Eq. (12) estimate the conditional score.

Theorem 2 (Treeffuser One-Dimensional Objectives). *Denote $U^* = (U_1^*, \dots, U_{d_y}^*)$. Then for almost all $\mathbf{x}, \mathbf{y}, t$ with respect to $\pi(\mathbf{x}, \mathbf{y})$ and the Lebesgue measure on $t \in [0, T]$, we have*

$$\nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y}) = \frac{U^*(\mathbf{y}, t, \mathbf{x})}{\sigma(t)}. \quad (13)$$

Each problem in Eq. (12) is a GBT problem where the notation within Eq. (4) corresponds to $F := U_k$, $\mathbf{a} := (h(\zeta, t, \mathbf{y}), t, \mathbf{x}) \in \mathbb{R}^{d_y+1+d_x}$, $b := -\zeta_k \in \mathbb{R}$ and L is the square loss. We note that the noise scaling reparametrization, $S = U/\sigma$, is key to stabilizing the learning process; see the ablation study in Appendix G.

Finally, Treeffuser approximates the expectations in Eq. (12) with Monte Carlo sampling. For each sample (\mathbf{x}, \mathbf{y}) from the dataset \mathcal{D} , Treeffuser samples R pairs of $(t, \zeta) \sim \text{Uniform}([0, 1]) \otimes \mathcal{N}(0, I_{d_y})$ and creates new datasets \mathcal{D}^k containing $R \cdot n$ datapoints of the form $((h(\zeta, t, \mathbf{y}), t, \mathbf{x}), -\zeta_k)$, one \mathcal{D}^k per dimension of \mathbf{y} . Then, each of these datasets is given to a standard GBT algorithm, such as LightGBM [22] or XGBoost [21]. Our implementation of Treeffuser uses LightGBM. Algorithm 1 details this procedure.

3.3 Sampling and Probabilistic Predictions

Treeffuser provides samples from the probabilistic predictive distribution $\pi(\mathbf{y} \mid \mathbf{x})$. It does so as in standard unconditional models by plugging in the GBT-estimated conditional score from Eq. (12)

into a numerical approximation of the SDE in Eq. (8). While Treeffuser is compatible with any SDE solver, our implementation leverages Euler-Maruyama [27] due to its good balance between accuracy and the number of function evaluations. In general, we found good performance with as few as 50 steps. Algorithm 2 implements this procedure.

The samples generated by Treeffuser can then be used to estimate means, quantiles, probability intervals, expectations, or any other quantity of interest that depends on the response distribution.

3.4 Limitations

The design of Treeffuser offers advantages in terms of usability, versatility, and robustness. But it also comes with a few limitations. First, the diffusion process is theoretically defined to only model continuous responses y . However, count and other forms of discrete responses are common in the probabilistic modeling of tabular data. While our experiments show that this limitation does not prevent Treeffuser from outperforming comparable methods on these kinds of data, there may be opportunities for further improvement with direct modeling of discrete outcomes. Second, Treeffuser does not offer a closed-form density and must solve an SDE to generate samples. This sampling process, in contrast with the fast training, can become expensive when many samples per datapoint x are required.

4 Related Work

Treeffuser builds on advances in diffusion models to form probabilistic predictions from tabular data.

Diffusion models. Diffusion models excel at learning complex unconditional distributions on a range of data, such as images [28–30], molecules [31], time series [32], and graphs [33]. A common task is conditional generation, where the goal is to generate samples from a distribution conditioned on features. There are two approaches to this objective. One approach is to use guidance methods by which the score of an unconditional diffusion model is altered during generation to mimic the score of the conditional distribution [11, 34, 35]. This approach is especially popular for inverse problems [11, 36, 37]. Another approach is to train a conditional model from the start, incorporating the conditioning information during training [24–26, 38, 39]. This is the approach adopted by Treeffuser.

Treeffuser contributes to a recent line of work that applies diffusions to tabular data. This includes deep learning approaches for data generation [40], probabilistic regression [41], and missing data imputation [42]. Among these methods, CARD [41] uses neural-net based diffusions for probabilistic predictions and thus is most similar to Treeffuser in scope. We attempted to include it in our experiments using the implementation from Lehmann [43], but returned very poor results. We therefore excluded it from our testbed. The imputation of missing data has been recently extended to gradient-boosted trees [44].

Probabilistic prediction for tabular data Probabilistic prediction for tabular data can be classified into parametric and non-parametric methods based on their assumptions about the likelihood shape. Parametric tree-based methods include NGBoost [14] and PBGM [45]. NGBoost uses natural gradients to optimize a scoring rule, while PBGM sequentially updates the mean and standard deviation for predictions. DRFs obtain maximum likelihood estimates by using this criteria to choose splits. Neural-based parametric methods include Bayesian Neural Networks [46], MC Dropout [47], and Deep Ensembles [17]. Notably these methods are all indirectly or directly Bayesian. Another approach, normalizing flows, transforms a latent distribution via an invertible neural network [48] and has been applied to tabular data [49]. Non-parametric methods are often tree-based, such as Quantile Regression Forests [16], Distributional Random Forests (DRF) [50], and IBUG [15]. Quantile Regression Forests approximate the inverse cumulative distribution function by minimizing pinball loss, while DRF and IBUG use a tree-based similarity metric to weight training data for predictions. These methods are baselines in our empirical studies, with detailed descriptions in Appendix F.1.

5 Empirical studies

We demonstrate Treeffuser across three settings: synthetic data, standard UCI datasets [51], and sales data from Walmart. We find that Treeffuser outperforms state-of-the-art methods [15, 17, 50];

Dataset	N, d_x, d_y	Deep ensembles	NGBoost (Gaussian)	iBUG (XGBoost)	Quantile regression	DRF	Treeffuser	Treeffuser (no tuning)	
bike	17,379,12,1	1.61 ± 0.05	7.15 ± 0.18	1.88 ± 0.11	1.85 ± 0.07	2.15 ± 0.06	1.60 ± 0.05	1.64 ± 0.05	$\times 10^1$
energy	768, 8, 2	5.00 ± 0.71	4.78 ± 0.49	NA	NA	5.43 ± 0.69	3.07 ± 0.40	3.32 ± 0.48	$\times 10^0$
kin8nm	8,192, 8, 1	3.59 ± 0.12	9.48 ± 0.29	7.74 ± 0.41	6.63 ± 0.16	9.44 ± 0.20	5.89 ± 0.14	5.88 ± 0.17	$\times 10^{-2}$
movies	7,415, 9, 1	2.94 ± 0.35	×	3.42 ± 0.52	7.90 ± 0.68	5.57 ± 0.61	2.68 ± 0.31	2.69 ± 0.28	$\times 10^7$
naval	11,934,17,1	4.11 ± 0.39	4.43 ± 0.19	3.20 ± 0.17	16.86 ± 2.46	4.55 ± 0.16	2.02 ± 0.08	2.46 ± 0.07	$\times 10^{-4}$
news	39,644,58,1	2.53 ± 0.27	×	3.89 ± 1.19	2.32 ± 0.17	1.98 ± 0.16	1.98 ± 0.17	1.98 ± 0.16	$\times 10^3$
power	9,568, 4, 1	2.06 ± 0.10	2.01 ± 0.13	1.61 ± 0.07	5.40 ± 0.12	1.90 ± 0.11	1.49 ± 0.07	1.52 ± 0.07	$\times 10^0$
superc.	21,263,81,1	4.89 ± 0.31	5.24 ± 0.50	4.14 ± 0.28	3.79 ± 0.14	4.32 ± 0.16	3.52 ± 0.13	3.60 ± 0.15	$\times 10^0$
wine	6,497, 12, 1	3.59 ± 0.10	3.82 ± 0.11	3.25 ± 0.14	3.24 ± 0.14	3.30 ± 0.12	2.59 ± 0.13	2.67 ± 0.13	$\times 10^{-1}$
yacht	308, 6, 1	4.86 ± 1.38	3.67 ± 1.47	3.75 ± 1.24	3.53 ± 1.30	7.73 ± 2.43	3.11 ± 0.99	3.39 ± 0.97	$\times 10^{-1}$

Table 1: CRPS (lower is better) by dataset and method. \times indicates the method failed to run, and NA that the method is not directly applicable to multivariate outputs. Standard deviations are measured with 10-fold cross-validation. For each dataset, the two best methods are bolded. Treeffuser provides the most accurate probabilistic predictions, even with default hyperparameters (no tuning). Powers of ten are factorized out of each row in the rightmost column.

it can capture complex distributions with multimodal, inflated, or multivariate responses; it achieves better probabilistic predictions on many real-world datasets and it produces more accurate sales forecasts on Walmart data. We also find that Treeffuser can outperform other (tuned) methods without dataset-specific hyperparameter tuning. We provide the code for Treeffuser at <https://github.com/blei-lab/treeffuser>.

5.1 Probabilistic prediction on synthetic data

We show the flexibility of Treeffuser on three probabilistic regression tasks that are difficult to model with standard parametric models. We also show its competitive performance on Gaussian data when compared to methods that posit a Gaussian likelihood.

Treeffuser vs. complex response functions. We design three difficult probabilistic regression tasks: a multimodal response where the number of modes changes with x , an inflated response with support that changes with x , and a multivariate response with nontrivial covariance structure. The data generation processes for these datasets are detailed in Appendix E. For each dataset, we generate 10,000 observations (x_i, y_i) and train Treeffuser with its default parameters (see Appendix C).

Fig. 1 shows the histograms of 1,000 Treeffuser samples for three different values of x and compares them with the true conditional densities of the response. Treeffuser captures the conditional distribution for all values of x across all settings. For multimodal data, Treeffuser correctly detects the modes of the distribution without knowing their number a priori. For inflated data, it recovers the peak at the inflation point and the x -dependent support of the response. For multivariate data, it recovers the complex covariance structures.

Treeffuser vs. parametric oracles. We simulate Gaussian responses with a linear response function and compare Treeffuser to other parametric methods that assume a Gaussian likelihood. These methods serve as oracles in these experiments as they are informed by the true functional form of the conditional response distribution. Results and further details are reported in Appendix E. Among all methods, Treeffuser consistently ranks second best, only behind the Deep Ensemble oracle.

5.2 Probabilistic prediction on real-world datasets

We compare Treeffuser with state-of-the-art methods for probabilistic predictions on standard UCI datasets [51]. A detailed description of the baseline models can be found in Appendix F.1.

Metrics. We evaluate probabilistic predictions with the continuous ranked probability score (CRPS) and standard accuracy metrics.

CRPS is defined as $\text{CRPS}(F, y) = \int_{-\infty}^{\infty} (F(y') - \mathbb{1}(y \leq y'))^2 dy'$, where F is the cumulative distribution function of the predicted distribution $p(y | x)$ and y is the true observed value [52]. CRPS is a proper scoring rule, in that its expected value is minimized under the true conditional distribution of the response [53]. CRPS is usually preferred over the log-likelihood, which is also

a proper scoring rule, but is sensitive to the estimation of the tail densities [54]. Also, CRPS can readily be estimated from samples of $p(y | \boldsymbol{x})$, which is our setting with the non-parametric methods we evaluate. We evaluate CRPS by generating 100 samples from $p(y | \boldsymbol{x})$ for each \boldsymbol{x} . For evaluating multivariate responses, we report the average marginal CRPS over each dimension.

We also measure the quality of point predictions for each model. This is the ability to predict conditional means $\mathbb{E}[y | \boldsymbol{x}]$. We approximate $\mathbb{E}[y | \boldsymbol{x}]$ using 50 samples and evaluate the accuracy using the mean absolute error (MAE) and the root mean squared error (RMSE).

Experimental setup. We performed 10-folds cross-validation. For each fold, we tuned the hyperparameters of the methods using Bayesian optimization for 25 iterations, using 20% of the current fold’s training data as a validation set. Additional Bayesian optimization’s iterations did not change the results. We detail the search space of hyperparameters for each method in Appendix F.2.

Results. Table 1 presents CRPS by dataset and method. Treeffuser consistently provides the most accurate predictions across datasets, as measured by CRPS. Notably, it outperforms other methods even when initialized with its default parameters. There is no consistent runner-up: among parametric methods, deep ensemble does overall better than NGBoost and IBUG; quantile regression does well on some datasets but underperforms in others.

We find that Treeffuser also returns the best point predictions of $\mathbb{E}[y | \boldsymbol{x}]$, as reported in Table 4. For comparison, we report the accuracy of point predictions from vanilla XGBoost and LightGBM in Table 6 in Appendix F.3. These methods do not provide probabilistic predictions but are tailored for point predictions. As expected, XGBoost and LightGBM outperform or tie with all the probabilistic methods. In particular, they often tie with Treeffuser, suggesting that Treeffuser provides probabilistic prediction without sacrificing average point predictions.

Finally, we conducted an ablation study to investigate the impact of the noise-scaling reparametrization of the score function on Treeffuser’s performance. As detailed in Appendix G, noise scaling is key to achieving top accuracy and stability.

5.3 Sales forecasting with long tails and count data

model metric	Deep ensembles	IBUG	NGBoost Poisson	Quantile regression	Treeffuser	Treeffuser (no tuning)	
CRPS	7.05	7.75	6.86	7.11	6.44	6.62	$\times 10^{-1}$
RMSE	2.03	2.16	2.33	2.88	2.09	2.09	$\times 10^0$
MAE	0.97	1.04	0.99	1.01	0.99	0.99	$\times 10^0$

Table 2: Walmart dataset metrics (lower is better). The evaluation is on the last 30 days of data. Treeffuser provides the best probabilistic predictions alongside NGBoost Poisson. Deep ensembles excel at point predictions. Powers of tens are factorized out of each row in rightmost column.

We further demonstrate the applicability of Treeffuser on a publicly available dataset [55] for sales forecasting under uncertainty. The goal is to forecast the number of units sold for a product given features such as its price, its type (e.g., food, cloths), and past sales.

This task is challenging due to zero inflation and long tails in the distribution of item sales. (E.g., umbrella sales are typically low but can spike during rainy weeks.) It is even more challenging for a diffusion model like Treeffuser, which is designed for continuous responses and not count data.

We use five years of sales data from ten Walmart stores (a large American retail chain). We randomly select 1,000 products, training on 112,000 data points from the first 1,862 days and evaluating 10,000 other data points for the remaining 30 days.

In addition to the previous baselines, we include NGBoost Poisson, a parametric model specifically designed for count data. We evaluate the predictions returned by each method in three ways.

CRPS and accuracy metrics. First, we compute the same evaluation metrics as in the previous experiments. We benchmark against the methods in the experiment and the methods in section 5.2.

The results are reported in Table 2. We find that Treeffuser again proves competitive, achieving a better CRPS than all methods and comparable MAE and RMSE.

Posterior predictive checks. Second, we perform held-out predictive checks [56], examining six statistics on the number of items sold: the total count of zero sales, the highest sales figure, and sales figures at the 0.99, 0.999, and 0.9999 quantiles (lower quantiles are well captured by all methods). We produce probabilistic predictions of these quantities by returning their empirical distributions as induced by the samples generated by the models. Fig. 2 compares the observed values against the probabilistic predictions. Treeffuser best captures the proportion of zeros and performs as well as NGBost-Poisson in modeling the behavior of the tails.

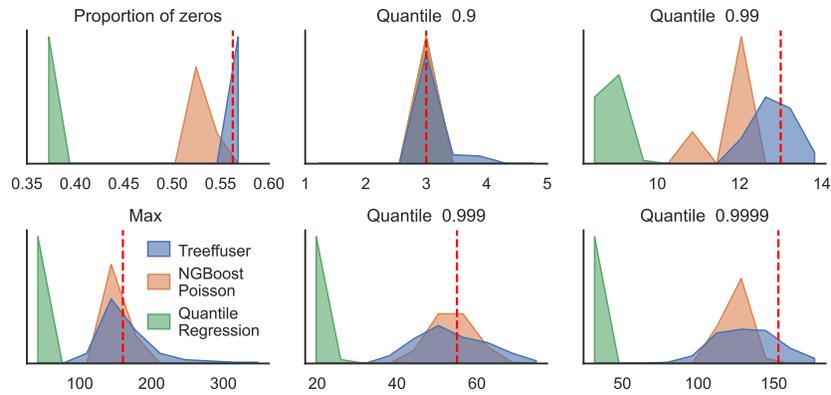


Figure 2: Posterior predictive checks for Treeffuser, NGBost Poisson, and quantile regression. Red dashed line shows the realized value on the test set. Treeffuser best captures the inflation point at zero and performs well on the tails.

Newsvendor model. Finally, we illustrate the practical relevance of accurate probabilistic predictions with an application to inventory management, using the newsvendor model [57]. Assume that every day we decide how many units q of an item to buy. We buy at a cost c and sell at a price p . However, the demand y is random, introducing uncertainty in our decision. The goal is to maximize the expected profit:

$$\max_q p \mathbb{E} [\min(q, y)] - cq.$$

The solution to the newsvendor problem is to buy $q = F^{-1} \left(\frac{p-c}{p} \right)$ units, where F^{-1} is the quantile function of the distribution of y .

We apply this model to evaluate the inventory decisions induced by each method on the Walmart dataset. To compute profits, we use the observed prices and assume a margin of 50% over all products. We let Treeffuser, NGBost-Poisson, and quantile regression learn the conditional distribution of the demand of each item, estimate their quantiles, and thus determine the optimal quantity to buy.

Fig. 3 plots the cumulative profits over the last 30 days of data. Treeffuser outperforms quantile regression by a large margin and performs comparably to NGBost-Poisson. This is coherent with our PPC results in Fig. 2, showing better quantile estimations for Treeffuser. This demonstrates that Treeffuser delivers competitive probabilistic predictions even for count data responses, a scenario it was not specifically designed to handle.

5.4 Runtime Performance Overview

We measure Treeffuser’s performance in terms of training and inference speed across different datasets. On the M5 dataset, using default parameters, Treeffuser completed training in 53.2 seconds on a MacBook Pro M3 Max and generated 10,000 samples (one per test data point) in 2.53 seconds. We conducted additional benchmarking experiments on both the UCI and M5 datasets, with results detailed in Appendix H. (Further discussion on the model’s time complexity is also provided in that appendix.) Details about the computational resources used in all of our experiments are available in Appendix D.

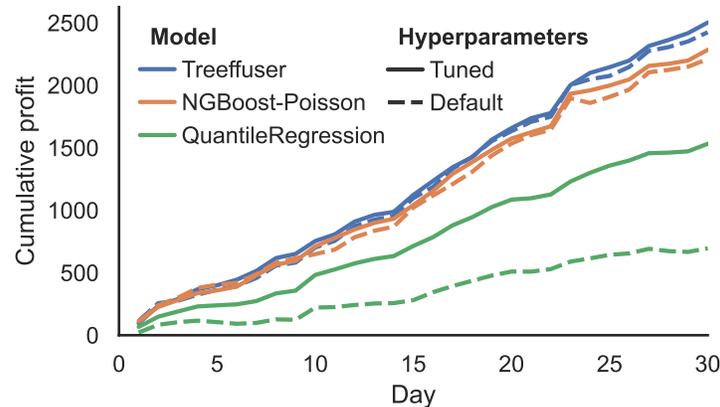


Figure 3: Cumulative profits by method on an inventory management problem. Treefuser produces more accurate probabilistic predictions yielding higher profits.

6 Conclusion

We have introduced Treefuser, a new model for probabilistic prediction from tabular data.

Treefuser combines conditional diffusions models with gradient-boosted trees. It can capture arbitrarily complex distributions without requiring any data-specific modeling or tuning. It is amenable to fast CPU learning and naturally handles categorical data and missing values. We have demonstrated that Treefuser outperforms state-of-the-art methods in probabilistic regression across datasets and metrics. These characteristics make Treefuser a flexible, easy-to-use, and robust model for probabilistic predictions.

One limitation of our diffusion-based approach is the need to numerically solve an SDE to generate samples, which can be costly when producing many samples. Recent advances, such as progressive distillation [58] and consistency models [59], address this issue. Applying these methods to Treefuser is a direction for future work.

References

- [1] Aristeidis Mystakidis, Evangelia Ntozi, Konstantinos Afentoulis, Paraskevas Koukaras, Paschalis Gkaidatzis, Dimosthenis Ioannidis, Christos Tjortjis, and Dimitrios Tzovaras. Energy generation forecasting: Elevating performance with machine and deep learning. *Computing*, 105(8):1623–1645, 2023.
- [2] Jillian M. Clements, Di Xu, Nooshin Yousefi, and Dmitry Efimov. Sequential deep learning for credit risk monitoring with tabular financial data, 2020. arXiv preprint arXiv:2012.15330.
- [3] Snigdha Somani, Adam J Russak, Filippo Richter, Simon Zhao, Akhil Vaid, Farah Chaudhry, Jason K De Freitas, Niyati Naik, Riccardo Miotto, Girish N Nadkarni, Jagat Narula, Edgar Argulian, and Benjamin S Glicksberg. Deep learning and the electrocardiogram: Review of the current state-of-the-art. *Europace*, 23(8):1179–1191, 2021.
- [4] Anshuman Gupta and Costas D Maranas. Managing demand uncertainty in supply chain planning. *Computers & Chemical Rngineering*, 27(8-9):1219–1227, 2003.
- [5] Yacov Y Haimes. *Risk Modeling, Assessment, and Management*. John Wiley & Sons, 2011.
- [6] Philippe Jorion. *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, 2007.
- [7] James W. Taylor and Kimberly S. Taylor. Combining probabilistic forecasts of covid-19 mortality in the united states. *European Journal of Operational Research*, 304(1):25–41, 2023.

- [8] Ziqiu Kang, Cagatay Catal, and Bedir Tekinerdogan. Machine learning applications in production lines: A systematic literature review. *Computers & Industrial Engineering*, 149:106773, 2020.
- [9] Anna L. Buczak and Erhan Guven. A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys and Tutorials*, 18(2): 1153–1176, 2016.
- [10] Thi-Thu-Huong Le, Yustus Eko Oktian, and Howon Kim. XGBoost for imbalanced multiclass classification-based industrial Internet of Things intrusion detection systems. *Sustainability*, 14 (14), 2022.
- [11] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020.
- [12] Jerome H Friedman. Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232, 2001.
- [13] Zhendong Wang, Yifan Jiang, Huangjie Zheng, Peihao Wang, Pengcheng He, Zhangyang Wang, Weizhu Chen, and Mingyuan Zhou. Patch Diffusion: Faster and more data-efficient training of diffusion models, 2023. arXiv preprint arXiv:2304.12526.
- [14] Tony Duan, Anand Avati, Daisy Yi Ding, Khanh K. Thai, Sanjay Basu, Andrew Y. Ng, and Alejandro Schuler. NGBoost: Natural gradient boosting for probabilistic prediction, 2020. arXiv preprint arXiv:1910.03225.
- [15] Jonathan Brophy and Daniel Lowd. Instance-based uncertainty estimation for gradient-boosted regression trees. In *Advances in Neural Information Processing Systems*, volume 35, pages 11145–11159, 2022.
- [16] Nicolai Meinshausen and Greg Ridgeway. Quantile regression forests. *Journal of Machine Learning Research*, 7(6), 2006.
- [17] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [18] Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [19] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.
- [20] Trevor Hastie, Robert Tibshirani, and Jerome H Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, volume 2. Springer, 2009.
- [21] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16. ACM, 2016.
- [22] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [23] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: Unbiased boosting with categorical features, 2019. arXiv preprint arXiv:1706.09516.
- [24] Georgios Batzolis, Jan Stanczuk, Carola-Bibiane Schönlieb, and Christian Etmann. Conditional image generation with score-based diffusion models, 2021. arXiv preprint arXiv:2111.13606.
- [25] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement, 2021. arXiv preprint arXiv:2104.07636.

- [26] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. In *Advances in Neural Information Processing Systems*, volume 34, pages 24804–24816. Curran Associates, Inc., 2021.
- [27] Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*. Stochastic Modelling and Applied Probability. Springer Berlin, Heidelberg, 1992.
- [28] Stability AI. Introducing stable diffusion. <https://stability.ai/blog/stable-diffusion-public-release>, 2022.
- [29] MidJourney. Midjourney. <https://www.midjourney.com/>, 2022.
- [30] OpenAI. Dall-e 2. <https://openai.com/index/dall-e-2/>, 2022.
- [31] Joseph L. Watson, David Juergens, Nathaniel R. Bennett, Brian L. Trippe, Jason Yim, Helen E. Eisenach, Woody Ahern, Andrew J. Borst, Robert J. Ragotte, Lukas F. Milles, Basile I. M. Wicky, Nikita Hanikel, Samuel J. Pellock, Alexis Courbet, William Sheffler, Jue Wang, Preetham Venkatesh, Isaac Sappington, Susana Vázquez Torres, Anna Lauko, Valentin De Bortoli, Emile Mathieu, Sergey Ovchinnikov, Regina Barzilay, Tommi S. Jaakkola, Frank DiMaio, Minkyung Baek, and David Baker. De novo design of protein structure and function with rfdiffusion. *Nature*, 620(7976):1089–1100, 2023.
- [32] Kashif Rasul, Calvin Seward, Ingmar Schuster, and Roland Vollgraf. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*, pages 8857–8868. PMLR, 2021.
- [33] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *International Conference on Artificial Intelligence and Statistics*, pages 4474–4484. PMLR, 2020.
- [34] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, pages 8780–8794, 2021.
- [35] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022. arXiv preprint arXiv:2207.12598.
- [36] Hyungjin Chung, Jeongsol Kim, Michael T. Mccann, Marc L. Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems, 2023. arXiv preprint arXiv:2209.14687.
- [37] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations, 2022. arXiv preprint arXiv:2108.01073.
- [38] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents, 2022. arXiv preprint arXiv:2204.06125.
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [40] Akim Kotelnikov, Dmitry Baranchuk, Ivan Rubachev, and Artem Babenko. TabDDPM: Modelling tabular data with diffusion models. In *International Conference on Machine Learning*, pages 17564–17579. PMLR, 2023.
- [41] Xizewen Han, Huangjie Zheng, and Mingyuan Zhou. CARD: Classification and regression diffusion models. In *Advances in Neural Information Processing Systems*, volume 35, pages 18100–18115, 2022.
- [42] Shuhan Zheng and Nontawat Charoenphakdee. Diffusion models for missing value imputation in tabular data, 2022. arXiv preprint arXiv:2210.17128.
- [43] Nils Lehmann. lightning-uq-box 0.1.0. <https://pypi.org/project/lightning-uq-box/>, 2024.

- [44] Alexia Jolicoeur-Martineau, Kilian Fatras, and Tal Kachman. Generating and imputing tabular data via diffusion and flow-based gradient-boosted trees. In *International Conference on Artificial Intelligence and Statistics*, pages 1288–1296. PMLR, 2024.
- [45] Olivier Sprangers, Sebastian Schelter, and Maarten de Rijke. Probabilistic gradient boosting machines for large-scale probabilistic regression. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '21. ACM, 2021.
- [46] Radford M Neal. *Bayesian Learning for Neural Networks*, volume 118. Springer Science & Business Media, 2012.
- [47] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International Conference on Machine Learning*, pages 1050–1059. PMLR, 2016.
- [48] Ivan Kobyzev, Simon J.D. Prince, and Marcus A. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979, 2021.
- [49] Pavel Izmailov, Polina Kirichenko, Marc Finzi, and Andrew Gordon Wilson. Semi-supervised learning with normalizing flows, 2019. arXiv preprint arXiv:1912.13025.
- [50] Domagoj Cevic, Loris Michel, Jeffrey Näf, Peter Bühlmann, and Nicolai Meinshausen. Distributional random forests: Heterogeneity adjustment and multivariate distributional regression. *Journal of Machine Learning Research*, 23(333):1–79, 2022.
- [51] Markelle Kelly, Rachel Longjohn, and Kolby Nottingham. The UCI machine learning repository, 2023. URL <https://archive.ics.uci.edu>.
- [52] James E Matheson and Robert L Winkler. Scoring rules for continuous probability distributions. *Management Science*, 22(10):1087–1096, 1976.
- [53] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [54] Mathias Blicher Bjerregård, Jan Kloppenborg Møller, and Henrik Madsen. An introduction to multivariate probabilistic forecast evaluation. *Energy and AI*, 4:100058, 2021.
- [55] Spyros Makridakis Addison Howard. M5 Forecasting - Accuracy, 2020. URL <https://kaggle.com/competitions/m5-forecasting-accuracy>.
- [56] Andrew Gelman, Xiao-Li Meng, and Hal Stern. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, 6:733–807, 1996.
- [57] Kenneth J Arrow, Theodore Harris, and Jacob Marschak. Optimal inventory policy. *Econometrica: Journal of the Econometric Society*, pages 250–272, 1951.
- [58] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models, 2022. arXiv preprint arXiv:2202.00512.
- [59] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models, 2023. arXiv preprint arXiv:2303.01469.
- [60] Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of Economic Perspectives*, 15(4):143–156, 2001.

Acronyms

CRPS continuous ranked probability score. 7, 8, 9

GBT gradient-boosted tree. 2, 3, 4, 5, 16

MAE mean absolute error. 8, 9

RMSE root mean squared error. 8, 9

SDE stochastic differential equation. 3, 4, 6, 14, 15, 16

A A primer on unconditional diffusion models

A diffusion model is a generative model that consists of two processes, a forward diffusion, and a reverse diffusion. The forward process takes an unknown target distribution $\pi(\mathbf{y})$ and continuously transforms it into a simple distribution $p_{\text{simple}}(\mathbf{y})$, typically a Gaussian. It does so by progressively adding noise to samples from π , such as our data.

The reverse process transforms the simple distribution back into the target distribution by denoising any perturbed samples from p_{simple} . If we can learn the reverse process, we can effectively access the target distribution π : first draw samples from p_{simple} and then use the reverse transformation, as guided by the score, to map them back to π .

One way to learn this noising-denoising process is to estimate the *score function*, the gradient of the log probability of the noisy data with respect to the data itself.

We detail the three main components of score-based diffusions—the forward process, the reverse process, and the estimation of the score function—below.

The forward diffusion process. A diffusion model is described by the following generative process:

$$\begin{cases} \text{Draw } \mathbf{y}(0) \sim \pi, \\ \text{Evolve } \mathbf{y}(t) \text{ until } T \text{ according to: } d\mathbf{y}(t) = f(\mathbf{y}(t), t)dt + g(t)d\mathbf{w}(t), \end{cases} \quad (14)$$

where the process of evolving the stochastic differential equation (SDE) involves a standard Brownian motion $\mathbf{w}(t)$ and model parameters $f : \mathbb{R}^{d_y} \times [0, T] \rightarrow \mathbb{R}^{d_y}$ and $g : [0, T] \rightarrow \mathbb{R}$, respectively called the *drift* and *diffusion* functions.

The SDE in Eq. (14) determines the evolution of $\mathbf{y}(t)$ over time, thereby inducing a distribution of $\mathbf{y}(t)$ at each t , that we write $p_t(\mathbf{y}(t))$. We further write the distribution of $\mathbf{y}(t)$ conditional on another $\mathbf{y}(u)$ as $p_{ut}(\mathbf{y}(t) | \mathbf{y}(u))$. In practice, f and g are simple functions such that $p_{0t}(\mathbf{y}(t) | \mathbf{y}(0))$ has a closed form. In fact, if the drift f is affine in $\mathbf{y}(t)$ and the diffusion function g is a scalar that does not depend on $\mathbf{y}(t)$, the distribution p_{0t} is always Gaussian, with a mean we write $m(t; \mathbf{y}(0))$ and a covariance multiple of the identity we write $\sigma(t)^2 I$.

We choose the time horizon T such that the resulting marginal distribution $p_T(\mathbf{y}(T)) = \int_{\mathbf{y}'} p_T(\mathbf{y}(T) | \mathbf{y}(0) = \mathbf{y}') \pi(\mathbf{y}') d\mathbf{y}'$ is a simple distribution p_{simple} agnostic to π . The intuition is that as the diffusion progresses, the initial density π is forgotten.

We detail our choice of f , g and T in Appendix C, which is the variance exploding SDE (VESDE) in Song et al. [11]. This choice induces a trivial closed form expression for the functions m and σ .

There are no model parameters to learn in the diffusion model, and the conditional distributions $p_{0t}(\mathbf{y}(t) | \mathbf{y}(0))$ are trivial to compute. The challenge is to learn how to compute (or sample from) the conditional distribution $p_{Tt}(\mathbf{y}(t) | \mathbf{y}(T))$. This is called *reversing* the diffusion process. Knowing p_{Tt} is key to learning the target distribution π , which can be written as $\pi(\mathbf{y}(0)) = \mathbb{E}_{p_T}[p_{T0}(\mathbf{y}(0) | \mathbf{y}(T))]$ where $p_T = p_{\text{simple}}$ is known.

The reverse diffusion process. Consider the following SDE that is reversed in time from T to 0,

$$\begin{cases} \tilde{\mathbf{y}}(T) \sim p_{\text{simple}}, \\ d\tilde{\mathbf{y}}_t = [f(\tilde{\mathbf{y}}(t), t) - g(t)^2 \nabla_{\tilde{\mathbf{y}}(t)} \log p_t(\tilde{\mathbf{y}}(t))]dt + g(t)d\tilde{\mathbf{w}}(t), \end{cases} \quad (15)$$

where $\tilde{\mathbf{w}}(t)$ is a standard Brownian with reversed time and p_t is the density of $\mathbf{y}(t)$ defined by the forward SDE (1). Similar to the forward case, the reverse SDE induces a distribution \tilde{p}_t for $\tilde{\mathbf{y}}(t)$ at each t . Anderson [18] shows that, if we denote by \tilde{p}_{Tt} the conditional distribution of $\tilde{\mathbf{y}}(t)$ given $\tilde{\mathbf{y}}(T)$, then $\tilde{p}_{Tt} = p_{Tt}$ for any time t . In other words, $\tilde{\mathbf{y}}(t) | \tilde{\mathbf{y}}(T) = \mathbf{a}$ and $\mathbf{y}(t) | \mathbf{y}(T) = \mathbf{a}$ have the same distribution for all \mathbf{a} . Since we designed the diffusion such that $p_T \approx p_{\text{simple}} = \tilde{p}_T$, we have in particular $\tilde{\mathbf{y}}(t) \sim \mathbf{y}(t)$.

SDE solvers let us sample from distributions that are solutions of an SDE; hence, we can obtain samples of $\pi(\mathbf{y})$ by solving Eq. (15). However, the function $\mathbf{y} \mapsto \nabla_{\mathbf{y}} \log p_t(\mathbf{y})$, called the *score*, is usually unknown and needs to be estimated.

Estimating the score. Vincent [19] shows that the score can be estimated from trajectories of the SDE as the minimizer of the following objective function:

$$[\mathbf{y} \mapsto \nabla_{\mathbf{y}} \log p_t(\mathbf{y})] = \arg \min_{s \in \mathcal{S}} \mathbb{E}_t \mathbb{E}_{\pi} \mathbb{E}_{p_{0t}} \left[\left\| \nabla_{\mathbf{y}(t)} \log p_{0t}(\mathbf{y}(t) \mid \mathbf{y}(0)) - s(\mathbf{y}(t), t) \right\|^2 \right], \quad (16)$$

where $\mathcal{S} = \{s : \mathbb{R}^{d_y} \times [0, T] \rightarrow \mathbb{R}^{d_y}\}$ is the set of all possible score functions indexed by time t , and \mathbb{E}_t can be an expectation over any distribution of t whose support is exactly $[0, T]$. Since p_{0t} is a normal distribution $\mathcal{N}(m(t; \mathbf{y}(0)), \sigma(t)^2 I)$, we have

$$\nabla_{\mathbf{y}(t)} \log p_{0t}(\mathbf{y}(t) \mid \mathbf{y}(0)) = \frac{m(t; \mathbf{y}(0)) - \mathbf{y}(t)}{\sigma(t)^2}. \quad (17)$$

In practice, the expectations in Eq. (16) are approximated using the empirical distribution formed by the observed \mathbf{y} for \mathbb{E}_{π} , and by sampling the simple and known forward SDE trajectories $\mathbf{y}(t) \mid \mathbf{y}(0) = \mathbf{y}$ for $\mathbb{E}_{p_{0t}}$. The objective can then be minimized by parametrizing s with any function approximator, e.g., a neural network. The score estimator can then be used to reverse the diffusion process and estimate the target distribution $\pi(\mathbf{y})$.

B Proofs of the main theorems

In this section, we provide the proofs for the two main theorems in the text.

Theorem 1 (Optimal Conditional Objective). *Define S^* as the solution of Eq. (9). Then, for almost all $\mathbf{x}, \mathbf{y}, t$ with respect to $\pi(\mathbf{x}, \mathbf{y})$ and the Lebesgue measure on $t \in [0, T]$, we have*

$$S^*(\mathbf{y}, t, \mathbf{x}) = \nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y}). \quad (10)$$

Proof. For conciseness, define for any function $s : \mathbb{R}^{d_y} \times [0, T] \rightarrow \mathbb{R}^{d_y}$ the quantity: $r(\mathbf{x}, s) = \mathbb{E}_t \mathbb{E}_{\pi(\mathbf{y}_{\mathbf{x}}(0))} \mathbb{E}_{p_{\mathbf{x}, 0t}} \left[\left\| \nabla_{\mathbf{y}_{\mathbf{x}}(t)} \log p_{0t}(\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(0)) - s(\mathbf{y}_{\mathbf{x}}(t), t, \mathbf{x}) \right\|^2 \right]$. Also write for any function $S : \mathbb{R}^{d_y} \times [0, T] \times \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$ and $\mathbf{x} \in \mathbb{R}^{d_x}$ the function $S_{\mathbf{x}} : (\mathbf{y}, t) \mapsto S(\mathbf{y}, t, \mathbf{x})$.

The function S^* is characterized as, $S^* \in \arg \min_{S \in \mathcal{S}^+} \mathbb{E}_{\pi(\mathbf{x})} [r(\mathbf{x}, S_{\mathbf{x}})]$ where $\mathcal{S}^+ = \{S : \mathbb{R}^{d_y} \times [0, T] \times \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}\}$

Define $\Omega = \{\mathbf{x} \in \mathbb{R}^{d_x} \mid S^*(\mathbf{y}, t, \mathbf{x}) \notin \arg \min_s r(\mathbf{x}, s)\}$, we will show that $\pi(\Omega) = 0$.

By [19], we know that $\nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y}) \in \arg \min_s r(\mathbf{x}, s)$ for all \mathbf{x} , so by definition of Ω , we have:

$$\forall \mathbf{x} \in \Omega, \quad r(\mathbf{x}, S_{\mathbf{x}}^*) > r(\mathbf{x}, (\mathbf{y}, t) \mapsto \nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y})). \quad (18)$$

We further have for any other $x \notin \Omega$, $r(\mathbf{x}, S_{\mathbf{x}}^*) = r(\mathbf{x}, (\mathbf{y}, t) \mapsto \nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y}))$.

If $\pi(\Omega) > 0$, then integrating Eq. (18) over $\pi(x)$ will yield $\mathbb{E}_{\pi(\mathbf{x})} [r(\mathbf{x}, S_{\mathbf{x}}^*)] > \mathbb{E}_{\pi(\mathbf{x})} [r(\mathbf{x}, ((\mathbf{y}, t, \mathbf{x}) \mapsto \nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y}))_{\mathbf{x}})]$, which is not possible by definition of S^* . Hence $\pi(\Omega) = 0$.

Hence we have $S_{\mathbf{x}}^* \in \arg \min_s r(\mathbf{x}, s)$ for almost any x , so by [19] again, we can conclude that $S^*(\mathbf{y}, t, \mathbf{x}) = \nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y})$ for almost any $\mathbf{x}, t, \mathbf{y}$. □

Theorem 2 (Treeffuser One-Dimensional Objectives). *Denote $U^* = (U_1^*, \dots, U_{d_y}^*)$. Then for almost all $\mathbf{x}, \mathbf{y}, t$ with respect to $\pi(\mathbf{x}, \mathbf{y})$ and the Lebesgue measure on $t \in [0, T]$, we have*

$$\nabla_{\mathbf{y}} \log p_{\mathbf{x}, t}(\mathbf{y}) = \frac{U^*(\mathbf{y}, t, \mathbf{x})}{\sigma(t)}. \quad (13)$$

Proof. By definition in Eq. (12), we have

$$U_k^* = \arg \min_{U_k} \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim \pi} \mathbb{E}_t \mathbb{E}_{\zeta \sim \mathcal{N}(0, I_{d_y})} \left[(\zeta_k - U_k(h(\zeta, t, \mathbf{y}), t, \mathbf{x}))^2 \right].$$

With Theorem 1, we have $S^*(\mathbf{y}, t, \mathbf{x}) = \nabla_{\mathbf{y}} \log p_{\mathbf{x},t}(\mathbf{y})$ almost everywhere, with S^* defined as

$$S^* \in \arg \min_{S \in S^+} \mathbb{E}_{\pi(\mathbf{x})} \mathbb{E}_t \mathbb{E}_{\pi(\mathbf{y}_{\mathbf{x}}(0))} \mathbb{E}_{p_{\mathbf{x},0t}} \left[\left\| \nabla_{\mathbf{y}_{\mathbf{x}}(t)} \log p_{0t}(\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(0)) - S(\mathbf{y}_{\mathbf{x}}(t), t, \mathbf{x}) \right\|^2 \right].$$

We have:

$$\mathbb{E}_{\pi(\mathbf{x})} \mathbb{E}_t \mathbb{E}_{\pi(\mathbf{y}_{\mathbf{x}}(0))} \mathbb{E}_{p_{\mathbf{x},0t}} \left[\left\| \nabla_{\mathbf{y}_{\mathbf{x}}(t)} \log p_{0t}(\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(0)) - S(\mathbf{y}_{\mathbf{x}}(t), t, \mathbf{x}) \right\|^2 \right] \quad (19)$$

$$= \mathbb{E}_{\pi(\mathbf{x})} \mathbb{E}_t \mathbb{E}_{\pi(\mathbf{y}_{\mathbf{x}}(0))} \mathbb{E}_{p_{\mathbf{x},0t}} \left[\left\| \frac{m(t; \mathbf{y}_{\mathbf{x}}(0)) - \mathbf{y}_{\mathbf{x}}(t)}{\sigma(t)^2} - S(\mathbf{y}_{\mathbf{x}}(t), t, \mathbf{x}) \right\|^2 \right] \quad (20)$$

$$= \sum_{k=1}^{d_y} \mathbb{E}_{\pi(\mathbf{x})} \mathbb{E}_t \mathbb{E}_{\pi(\mathbf{y}_{\mathbf{x}}(0))} \mathbb{E}_{p_{\mathbf{x},0t}} \left[\frac{1}{\sigma(t)} \left(\frac{m(t; \mathbf{y}_{\mathbf{x}}(0))_k - \mathbf{y}_{\mathbf{x}}(t)_k}{\sigma(t)} - \sigma(t) S(\mathbf{y}_{\mathbf{x}}(t), t, \mathbf{x})_k \right) \right]^2 \quad (21)$$

$$= \sum_{k=1}^{d_y} \mathbb{E}_{\pi(\mathbf{x}, \mathbf{y})} \mathbb{E}_t \mathbb{E}_{\zeta \sim \mathcal{N}(0, I_{d_y})} \left[\frac{1}{\sigma(t)} \left(-\zeta_k - \sigma(t) S(\mathbf{y} + \sigma(t)\zeta, t, \mathbf{x})_k \right) \right]^2 \quad (22)$$

where we used the following facts:

- from Eq. (19) to Eq. (20): the closed-form expression of the score S from Eq. (17),
- from Eq. (20) to Eq. (21): expanding the norm and switching the expectations with the finite sum over k ,
- from Eq. (21) to Eq. (22): reparametrizing the expectation of $\mathbf{y}_{\mathbf{x}}(t) \mid \mathbf{y}_{\mathbf{x}}(0)$ which by definition is a normal distribution with mean $m(t; \mathbf{y}_{\mathbf{x}}(0))$ and variance $\sigma(t)^2$.

The final manipulation is to remember that theorem 1 and the theorems from Vincent [19] are valid with expectations \mathbb{E}_t against any strictly positive measure of t over $[0, T]$. In particular, we can absorb $\frac{1}{\sigma(t)}$ as a reweighted non-negative measure, and we obtain exactly the definition of U_k^* by defining $U(\mathbf{y}, t, \mathbf{x}) = \sigma(t) S(\mathbf{y} + \sigma(t)\zeta, t, \mathbf{x})$ which concludes the proof. \square

C Treffuser default configuration

We use the following configuration as defaults for Treffuser.

Forward diffusion. We use the variance exploding SDE [11] for all experiments, defined by setting:

$$f(\mathbf{y}, t) = 0 \quad \text{and} \quad g(t) = \sqrt{\frac{d[\sigma(t)^2]}{dt}}, \quad (23)$$

where σ is a given increasing function defined by,

$$\sigma(t) = \alpha_{\min} \left(\frac{\alpha_{\max}}{\alpha_{\min}} \right)^t \quad (24)$$

and $\alpha_{\min} = 0.01$ and $\alpha_{\max} = 20$. For all our experiments we let $t \in [0, 1]$ (i.e. $T = 1$).

Gradient-boosted tree (GBT) parameters and dataset repetitions. We provide a short description of each hyper-parameter of the model alongside the default value. Treffuser uses LightGBM [22] to learn the GBTs.

- `n_estimators` (3000): Specifies the maximum number of trees that will be fit, regardless of whether the stopping criterion is met.
- `learning_rate` (0.1): Specifies the shrinkage to use for every tree.
- `num_leaves` (31): Specifies the maximum number of leaves a tree can have.
- `early_stopping_rounds` (50): Specifies how long to wait without a validation loss improvement before stopping.
- `n_repeats` (30): Specifies how many Monte Carlo samples to draw per data point to estimate $\mathbb{E}_{t,\zeta}$ in equation Eq. (9).

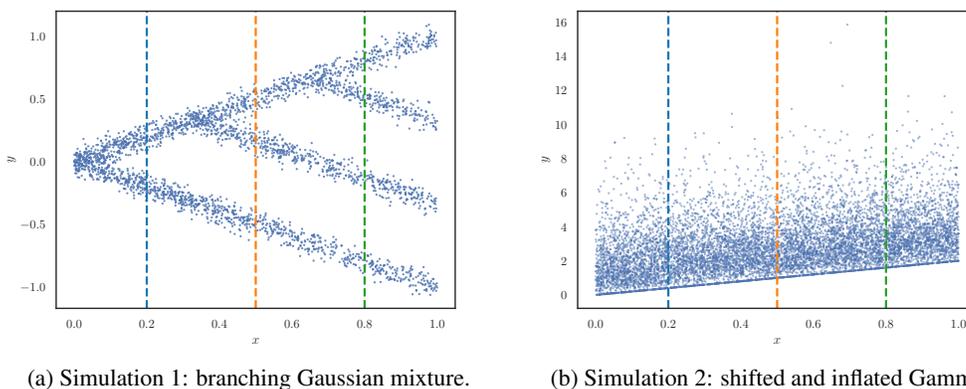


Figure 4: Visualization of ground-truth samples for the one-dimensional synthetic datasets used in the empirical studies.

D Experiments: description of computer resources

All benchmark tasks presented in Section 5.1 and Section 5.2 were run on a cluster with one job per triplet of (model, dataset, split index).

The real world experiments totalled 1135h for 700 tasks (10 datasets, 7 methods, 10 splits). Each task was allocated 4 cpus.

The synthetic experiments totalled 290 hours, for 800 tasks (16 datasets, 5 methods, 10 splits). Each task was allocated 2 CPU.

The M5 experiments were run on a single MacBook Pro over a couple hours. Other figures such as Fig. 1 were also generated on a MacBook pro in a few seconds.

E Experiments on synthetic data (supplement)

E.1 Arbitrarily complex synthetic data

We provide details on the three synthetic data experiments presented in Section 5.1 and illustrated in Fig. 1. Each experiment introduces a different distribution of the response variable y given the covariate x . The first experiment generates multimodal responses from a branching Gaussian mixture, the second experiment generates inflated responses from a mixture of a shifted Gamma distribution and an atomic measure, and the third experiment generates two-dimensional responses from a nonlinear multioutput regression.

We provide additional visualization of samples for the one-dimensional datasets in Fig. 4.

Models We assume that $x \sim \text{Uniform}([0, 1])$ and model the conditional distribution of the response y given x differently for each of the synthetic dataset. Scatter plots of synthetic data from these distributions are provided in Fig. 4

Branching Gaussian mixture. We generate multimodal responses from a mixture of equally-weighted Gaussians. In our experiments, we fix the scale $\sigma = 0.05$ and let the number of components and their means scale with x as follows:

$$y \mid x \sim \begin{cases} \text{GaussianMixture}_\sigma(x, -x), & 0 \leq x \leq 1/3; \\ \text{GaussianMixture}_\sigma(x, 2/3 - x, -x), & 1/3 < x \leq 2/3; \\ \text{GaussianMixture}_\sigma(x, 4/3 - x, 2/3 - x, -x), & 2/3 \leq x \leq 1; \end{cases} \quad (25)$$

where $\text{GaussianMixture}_\sigma(\mu_1, \mu_2, \dots, \mu_K)$ denotes a Gaussian mixture distribution with K equally-weighted components, each with mean μ_k and scale σ .

Covariate inflated Gamma. We generate inflated responses through the following mixture of a shifted Gamma distribution and an atomic measure:

$$z \sim \text{Bernoulli}(p), \quad y \mid x, z \sim \begin{cases} \text{ShiftedGamma}(k, \theta, x), & z = 0; \\ \delta_x, & z = 1; \end{cases} \quad (26)$$

where $\text{ShiftedGamma}(k, \theta, x)$ denotes a Gamma distribution with scale k and shape θ shifted by x . In words, given x , we set y to be equal to x with probability p , and to be drawn from a Gamma shifted by x otherwise. Hence, the conditional distribution of y given x is x -inflated and its support $[x, \infty)$ changes with x . In our experiments, we set $p = 0.15$, $k = 2$, $\theta = 1$.

Nonlinear multioutput regression. We generate two-dimensional responses from a density $p(\mathbf{y} \mid x) \propto \exp\left(-\frac{\text{dist}(\mathbf{y}, C(x))^2}{2\sigma^2}\right)$ where $\sigma = 0.05$, $\text{dist}(a, B)$ measures the shortest Euclidean distance from $a \in \mathbb{R}$ to the set $B \subset \mathbb{R}^2$, and $C(x)$ is a circular arc centered at $(0, 0)$ with radius $\text{radius}(x)$ and angles of extremities θ_1, θ_2 evolving as functions of x . Intuitively, they interpolate the circular arc represented in Section 1.

The functions are defined as follows.

- If $x \leq 0.5$
 - $\text{radius}(x) = \ell(x) \cdot 1 + r(x) \cdot 0.1$,
 - $\theta_0(x) = \ell(x) \cdot (-0.05) + r(x) \cdot (-0.375)$,
 - $\theta_1(x) = \ell(x) \cdot 0.3 + r(x) \cdot 0.625$,

where $\ell(x) = (0.5 - x)/(0.5 - 0.17)$, $r(x) = 1 - \ell(x)$.
- If $x > 0.5$
 - $\text{radius}(x) = \ell(x) \cdot 0.1 + r(x) \cdot 1$,
 - $\theta_0(x) = \ell(x) \cdot 0.125 + r(x) \cdot 0.45$,
 - $\theta_1(x) = \ell(x) \cdot 1.125 + r(x) \cdot 0.8$,

where $\ell(x) = (0.83 - x)/(0.83 - 0.5)$, $r(x) = 1 - \ell(x)$.

E.2 Data with normal likelihood

We study the performance of Treefuser in a setting where a simpler model is better adjusted to the generating process. In particular, we let N be number of data points sampled and d_x the dimension and use a linear generative model

$$\begin{aligned} \beta &\sim \mathcal{N}(0, I_{d_x}) \\ \mathbf{x}_i &\sim \mathcal{N}(0, I_{d_x}) \quad i \in [N] \\ \mathbf{y}_i &\sim \mathcal{N}(10\beta^\top \mathbf{x}, \sigma^2) \quad i \in [N] \end{aligned}$$

where $\sigma^2 = 1$. The results are provided in table 3 using the same configuration described in section section 5.2.

Results From Table 3, we observe that while Treefuser remains competitive, Deep Ensembles and NGBoost now demonstrate either equal or very close performance. This is expected since both models assume a normal distribution for the outcomes and the model assumed by Deep Ensembles closely resembles the real generative model. Interestingly, Treefuser outperforms both iBUG and quantile regression (QReg), suggesting that its superior performance is not merely due to using LightGBM or XGBoost. Finally, we note that Treefuser continues to perform well in low-data situations compared to other models.

F Experiments on standard datasets (supplement)

F.1 Baseline methods.

We briefly describe the methods used as baselines in our empirical studies.

N	d_x	Deep Ens. (oracle)	IBUG (oracle)	NGBoost (oracle)	QReg	Treeffuser
100	1	2.62 ±1.67	8.65±10.30	4.08±4.35	5.00±4.59	4.01 ±4.11 ×10 ⁻²
100	5	3.72 ±1.60	4.45±1.67	4.63±1.75	4.69±2.02	3.86 ±1.67 ×10 ⁻¹
100	10	4.04 ±1.77	5.31±2.23	4.58 ±1.60	5.13±2.20	4.90±2.08 ×10 ⁻¹
100	20	4.27 ±2.65	5.15±2.29	4.49 ±2.30	5.87±3.03	4.63±2.19 ×10 ⁻¹
500	1	0.84 ±0.41	1.58±0.66	0.87 ±0.41	1.36±0.73	1.11±0.44 ×10 ⁻²
500	5	3.19 ±0.27	4.00±0.55	3.82 ±0.58	4.44±0.74	3.85±0.51 ×10 ⁻¹
500	10	3.44 ±0.71	4.25±0.87	4.03±0.90	4.51±1.10	3.98 ±0.83 ×10 ⁻¹
500	20	3.69 ±0.81	4.31 ±1.04	4.82±1.50	4.80±1.18	4.32 ±0.91 ×10 ⁻¹
1000	1	2.51 ±1.08	10.59±2.57	7.10 ±2.09	11.85±2.34	9.72±2.47 ×10 ⁻³
1000	5	3.18 ±0.26	3.67±0.39	3.73±0.33	3.66±0.34	3.52 ±0.21 ×10 ⁻¹
1000	10	3.29 ±0.42	4.01±0.63	4.01±0.54	4.23±0.76	3.95 ±0.59 ×10 ⁻¹
1000	20	3.40 ±0.29	4.25±0.34	4.11±0.47	4.27±0.49	4.00 ±0.38 ×10 ⁻¹
5000	1	0.56 ±0.25	9.25±2.71	3.87 ±2.26	10.36±2.96	10.00±2.49 ×10 ⁻³
5000	5	3.27 ±0.26	3.69±0.36	3.63±0.33	3.63±0.34	3.63 ±0.29 ×10 ⁻¹
5000	10	3.18 ±0.18	3.82±0.26	3.66±0.25	3.65±0.27	3.59 ±0.25 ×10 ⁻¹
5000	20	3.21 ±0.20	3.86±0.23	3.90±0.25	3.80±0.24	3.69 ±0.22 ×10 ⁻¹

Table 3: CRPS (lower is better) by dataset and method. Treeffuser provides comparable predictions to methods with parametric assumptions that match the generating process. The standard deviations are measured with 10-fold cross-validation.

NGBoost (parametric, tree-based). NGBoost [14] models the conditional distribution of the target variable using a parametric family of distributions whose parameters are predicted by a gradient-boosting algorithm. It uses natural gradients for more stable, accurate, and faster learning. In our experiments, we set the parametric model to be Gaussian.

Deep ensembles (parametric, nnet-based). A Deep ensemble [17] is a collection of neural networks that are individually trained to model a parametric conditional distribution $p(y|x)$. The neural networks are then combined to obtain a mixture. We set the parametric model to be Gaussian.

Quantile regression (nonparametric, tree-based). Quantile regression [60] estimates the quantiles of the conditional distribution $p(y|x)$. We implemented a GBT-based version of the method by fitting trees with inputs x and q , where q is the probability of the desired quantile. During training, we optimized the objective $\mathbb{E}_q \mathbb{E}_{x,y} [L(Q(x, y), q)]$, where L is the pinball loss and q is sampled from a zero-one uniform distribution. In our experiments, we used LightGBM as a GBT method.

IBUG (parametric, tree-based). IBUG [15] extends any GBT into a probabilistic estimator. Given an input instance x , it outputs a distribution around the prediction using the k -nearest training data points. The distance between a training instance x_0 and x depends on the number of co-occurrences of x_0 and x across the leaves of the ensemble. In our experiments, we used XGBoost and a Gaussian likelihood, following the default specifications of the method.

DRF (nonparametric, tree-based). DRF [50] grows a random forest by maximizing the differences in the response distributions across split groups with a distributional metric. Repeated randomization induces a weighting function that assesses how relevant a training data point is to a given input x . These weights are used to return a weighted empirical distribution of the response.

F.2 Hyperparameters of the methods.

Unless specified otherwise, we tuned the hyperparameters of each method using Bayesian optimization for 25 iterations. We used the following search space for each method (where $\llbracket a, b \rrbracket$ denotes the set of integers from a to b):

- Treeffuser
 - n estimators $\in \llbracket 100, 3000 \rrbracket$
 - n repeats $\in \llbracket 10, 50 \rrbracket$
 - learning rate $\in [0.01, 1]$ (log-uniform)
 - early stopping rounds $\in \llbracket 10, 100 \rrbracket$
 - num leaves $\in \llbracket 10, 100 \rrbracket$

- Quantile regression (QReg)
 - n estimators \in $\llbracket 100, 3000 \rrbracket$ (log-uniform)
 - n repeats \in $\llbracket 10, 100 \rrbracket$
 - learning rate \in $[0.01, 1]$
 - early stopping rounds \in $\llbracket 10, 100 \rrbracket$
 - num leaves \in $\llbracket 10, 100 \rrbracket$
- IBUG
 - k \in $\llbracket 20, 250 \rrbracket$
 - n estimators \in $\llbracket 10, 1000 \rrbracket$
 - learning rate \in $[0.01, 0.5]$ (log-uniform)
 - max depth \in $\llbracket 1, 100 \rrbracket$
- DRF
 - min node size \in $\llbracket 5, 30 \rrbracket$
 - num trees \in $\llbracket 250, 3000 \rrbracket$
- NGBOOST
 - n estimators \in $\llbracket 100, 10000 \rrbracket$
 - learning rate \in $[0.005, 0.2]$
- Deep ensemble
 - n layers \in $\llbracket 1, 5 \rrbracket$
 - hidden size \in $\llbracket 10, 500 \rrbracket$
 - learning rate \in $[10^{-5}, 10^{-2}]$ (log-uniform)
 - n ensembles \in $\llbracket 2, 10 \rrbracket$

For the methods that only return point predictions, we used the following search spaces:

- XGBoost
 - n estimators \in $\llbracket 10, 1000 \rrbracket$
 - learning rate \in $[0.01, 0.5]$ (log-uniform)
 - max depth \in $\llbracket 1, 100 \rrbracket$
- LightGBM
 - n estimators \in $\llbracket 10, 1000 \rrbracket$
 - learning rate \in $[0.01, 0.5]$ (log-uniform)
 - num leaves \in $\llbracket 10, 100 \rrbracket$
 - early stopping rounds \in $\llbracket 10, 100 \rrbracket$

F.3 Accuracy and calibration results.

dataset	N, d_x, d_y	Deep ensembles	NGBoost (Gaussian)	iBUG (XGBoost)	Quantile regression	DRF	Treefuser	Treefuser (no tuning)	
bike	17379,12,1	3.70±0.13	11.52±0.30	4.11±0.19	4.63±0.21	4.86±0.18	3.69±0.14	3.81±0.15	$\times 10^1$
energy	768, 8, 2	11.34±2.08	11.41±1.25	NA	NA	13.73±1.69	8.32±1.35	8.79±1.46	$\times 10^{-1}$
kin8nm	8192, 8, 1	0.64±0.02	1.71±0.06	1.38±0.08	1.18±0.04	1.72±0.04	1.06±0.03	1.06±0.02	$\times 10^{-1}$
movies	7415, 9, 1	9.35±1.90	×	9.83±1.44	18.23±2.61	15.87±2.59	8.85±1.66	8.63±1.38	$\times 10^7$
naval	11934,17,1	10.42±1.69	8.66±0.61	6.46±1.07	148.80±35.50	11.49±0.98	4.40±0.26	4.92±0.57	$\times 10^{-4}$
news	39644,58,1	1.95±2.66	×	1.18±0.40	1.12±0.40	1.08±0.41	1.10±0.40	1.09±0.41	$\times 10^4$
power	9568, 4, 1	3.81±0.32	3.66±0.35	3.06±0.32	36.27±0.56	3.65±0.35	2.93±0.30	3.02±0.30	$\times 10^0$
superc.	21263,81,1	11.26±0.74	11.25±0.87	9.67±0.54	9.36±0.48	10.79±0.56	9.08±0.48	9.21±0.53	$\times 10^0$
wine	6497, 12, 1	6.54±0.17	6.85±0.18	6.02±0.26	6.28±0.29	6.84±0.21	5.88±0.17	5.98±0.19	$\times 10^{-1}$
yacht	308, 6, 1	1.06±0.46	0.83±0.31	0.83±0.31	1.22±0.74	2.30±1.07	2.10±1.08	2.08±0.45	$\times 10^{-1}$

Table 4: RMSE (lower is better) by dataset and method. × indicates the method failed to run, and NA that the method is not directly applicable to multivariate outputs. The standard deviations are measured with 10-fold cross-validation. For each dataset, the two best methods are bolded. Treefuser provides the most accurate probabilistic predictions, even when initialized with defaults.

dataset	N, d_x, d_y	Deep ensembles	NGBoost (Gaussian)	iBUG (XGBoost)	Quantile regression	DRF	Treefuser	Treefuser (no tuning)	
bike	17379,12,1	3.28±0.97	17.79±0.68	5.77±1.04	2.55±0.29	3.45±0.43	1.73±0.62	1.86±0.87	$\times 10^{-2}$
energy	768, 8, 2	8.08±1.41	5.82±1.74	NA	NA	4.54±0.74	4.11±1.14	5.53±1.42	$\times 10^{-2}$
kin8nm	8192, 8, 1	2.58±1.50	4.70±0.44	2.99±0.58	8.47±0.80	2.79±0.52	3.85±0.96	3.34±1.22	$\times 10^{-2}$
movies	7415, 9, 1	5.75±1.38	×	7.56±3.31	49.93±0.07	1.31±0.47	2.71±1.36	5.17±1.32	$\times 10^{-2}$
naval	11934,17,1	13.31±1.90	5.50±0.75	3.86±2.59	15.05±0.58	17.37±0.40	7.75±0.41	9.10±0.39	$\times 10^{-2}$
news	39644,58,1	11.36±0.63	×	14.43±1.82	18.96±0.42	1.14±0.25	3.50±0.70	4.87±0.28	$\times 10^{-2}$
power	9568, 4, 1	2.52±1.18	2.53±1.00	3.67±1.42	4.75±0.93	2.37±0.58	4.36±1.31	1.77±0.74	$\times 10^{-2}$
superc.	21263,81,1	3.19±0.63	2.75±0.69	5.39±1.97	5.47±0.44	3.87±0.32	1.34±0.10	1.33±0.16	$\times 10^{-2}$
wine	6497, 12, 1	1.92±0.56	2.68±0.55	4.72±0.83	4.36±0.91	22.02±1.13	5.70±0.55	5.22±0.69	$\times 10^{-2}$
yacht	308, 6, 1	13.35±2.17	9.04±2.41	9.06±3.26	8.26±2.80	7.15±2.86	7.27±2.58	10.02±2.29	$\times 10^{-2}$

Table 5: MACE (lower is better) by dataset and method. × indicates the method failed to run, and NA that the method is not directly applicable to multivariate outputs. The standard deviations are measured with 10-fold cross-validation. For each dataset, the two best methods are bolded. Treefuser has a competitive calibration error across most datasets.

dataset	N, d_x, d_y	iBUG (XGBoost)	XGBoost	LightGBM	Treefuser	Treefuser (no tuning)	
bike	17379,12,1	4.11 ± 0.19	3.72 ± 0.13	3.79 ± 0.12	3.69 ± 0.14	3.81 ± 0.15	$\times 10^1$
kin8nm	8192, 8, 1	1.38 ± 0.08	1.12 ± 0.04	1.04 ± 0.03	1.06 ± 0.03	1.06 ± 0.02	$\times 10^{-1}$
movies	7415, 9, 1	9.83 ± 1.44	8.96 ± 1.60	9.08 ± 1.76	8.85 ± 1.66	8.63 ± 1.38	$\times 10^7$
naval	11934,17,1	6.46 ± 1.07	5.11 ± 0.45	6.68 ± 1.53	4.40 ± 0.26	4.92 ± 0.57	$\times 10^{-4}$
news	39644,58,1	1.18 ± 0.40	1.09 ± 0.40	1.09 ± 0.40	1.10 ± 0.40	1.09 ± 0.41	$\times 10^4$
power	9568, 4, 1	3.06 ± 0.32	2.97 ± 0.33	2.99 ± 0.32	2.93 ± 0.30	3.02 ± 0.30	$\times 10^0$
superc.	21263,81,1	9.67 ± 0.54	9.04 ± 0.48	9.21 ± 0.44	9.08 ± 0.48	9.21 ± 0.53	$\times 10^0$
wine	6497, 12, 1	6.02 ± 0.26	5.95 ± 0.23	5.90 ± 0.27	5.88 ± 0.17	5.98 ± 0.19	$\times 10^{-1}$
yacht	308, 6, 1	0.83 ± 0.31	2.19 ± 1.82	0.67 ± 0.40	2.10 ± 1.08	2.08 ± 0.45	$\times 10^{-1}$

Table 6: RMSE (lower is better) by dataset and method. This table extends Table 4 by includes vanilla XGBoost and LightGBM, alongside iBUG and Treefuser. The standard deviations are measured with 10-fold cross-validation. For each dataset, the two best methods are bolded.

G Ablation study on noise scaling

We evaluate the impact of including noise scaling in the score parametrization of Eq. (13), specifically $S(\mathbf{y}, t, \mathbf{x}) = U(\mathbf{y}, t, \mathbf{x})/\sigma(t)$, where $S(\mathbf{y}, t, \mathbf{x})$ denotes the score function of $\mathbf{y}(t) \mid \mathbf{x}$, U the gradient-boostered trees, and $\sigma(t)$ the noise schedule of the diffusion process. We compare Treefuser with and without noise scaling under the same experimental setup of Section 5.2.

Table 7 presents the CRPS by dataset and method. Treeffuser without noise scaling fails to achieve any reasonable CRPS. We believe Treeffuser without noise scaling may require extra model capacity (e.g., deeper trees and more training iterations) to handle the varying noise levels adequately. However, we did not further optimize it, as noise scaling provides state-of-the-art results without the need for additional complexity.

Treeffuser	bike	energy	kin8nm	naval	news	power	superc.	wine	yacht
with scaling	1.60±0.05	3.07±0.40	5.89±0.14	2.02±0.08	1.98±0.17	1.49±0.07	3.52±0.13	2.59±0.13	3.11±0.99
without scaling	176±211	5258±9667	204±240	154±175	62.1±78.4	289±296	112±162	30±53	18682±41026
	×10 ¹	×10 ⁻¹	×10 ⁻²	×10 ⁻⁴	×10 ³	×10 ⁰	×10 ⁰	×10 ⁻¹	×10 ⁻¹

Table 7: CRPS (lower is better) by dataset for Treeffuser with and without the noise scaling reparametrization in Eq. (13). Without scaling, Treeffuser does not produce meaningful results.

H Runtime and Complexity

	m5	movies	bike	energy	kin8nm	naval	news	power	superc.	wine	yacht
$\frac{\text{Time}}{\text{sample}}$	1.33	0.68	1.33	0.55	1.23	1.46	1.31	1.15	0.85	0.56	0.28

Table 8: Average running time in milliseconds for producing a sample from the model.

To demonstrate the efficiency of Treeffuser, we present the results of the following experiments and benchmarks:

1. We run Treeffuser on subsets of the M5 dataset (section 5.3) of various sizes and report the training time in Fig. 5.
2. We report the training time for running Treeffuser on the datasets used in the paper in Fig. 6.
3. We report in Table 8 the average runtime for generating a single sample after training Treeffuser. The average is calculated by sampling five thousand points after training the model with default settings and 50 discretization steps.

We conducted all of these experiments on a 2020 MacBook Pro with a 2.6 GHz 6-Core Intel Core i7 processor.

From these results we find that:

- Treeffuser is very fast at fitting moderately sized datasets, with runtime increasing linearly with dataset size.
- Sampling a single points is very fast ($\approx 10^{-3}$ seconds) yet, drawing many samples can become significant, e.g., with a large test set.

With respect to the time complexity of the model we note the following. If the time complexity of fitting a single GBT is of order $O(F(|\mathcal{D}|))$, where $F : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and $|\mathcal{D}|$ is the size of the dataset, then the time complexity of fitting Treeffuser is $O(d \times F(\text{n_repeats} \times |\mathcal{D}|))$, where d is the dimension of \mathbf{y} and n_repeats is the parameter that determines how many noisy versions of the dataset are sampled. Similarly, assuming constant time per GBT evaluation, the complexity of sampling is $O(\text{n_samples} \times \text{n_discretization_steps})$, where n_samples is the number of samples drawn from the model, and $\text{n_discretization_steps}$ is the number of function evaluations used by the SDE solver.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We propose a method for probabilistic prediction from tabular data. In our abstract and introduction, we make two central claims:

- Treeffuser is accurate, in that it outperforms state-of-the-art methods in probabilistic prediction from tabular data on standard UCI datasets. This is discussed in detail in Section 5.2 and Appendix F. Table 1, Table 4, and Table 5 present the results.
- Treeffuser is flexible, in that it adapts to complex distributions. This is shown in Fig. 1, which presents the results of our synthetic experiments, and in Section 5.1 and Appendix E, which provide the details.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitations of our method in Section 3.4. Treeffuser has two limitations: the sampling process can become expensive when many samples are required for a given input instance x ; the diffusion model is specifically designed to generate only continuous responses, which may not be suitable for count data. However, in Section 5.3, we demonstrate that Treeffuser outperforms other methods even on this kind of data.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We state two theorems, Theorem 1 and Theorem 2, and discuss their implications in Section 3.1 and Section 3.2, respectively. Proofs are given in Appendix B.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The details needed to reproduce our experiments are stated in Section 5 and Appendices C, E and F. We also provide an implementation of Treeffuser in the supplemental material. The notebooks with the experiments can be found in `testbed/notebooks/paper_notebooks`.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.

- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide an implementation of Treeffuser in the supplemental material. The notebooks with the experiments can be found in `testbed/notebooks/paper_notebooks`. For our experiments on real data, we use publicly available datasets and point to their references [51, 55].

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The details needed to reproduce our experiments are stated in Section 5 and Appendices C, E and F.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All the accuracy and calibration metrics reported in our tables present standard deviations computed with 10-fold cross validation.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Yes, we provide details in Appendix D: the number of tasks, the number of cpus and the total time taken. The memory is not reported as our work has no particular needs for memory and no particular configuration was necessary.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We provide full transparency in our methodologies and support reproducibility through shared code and datasets. There are no conflicts of interest to declare.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We are not releasing pre-trained models or datasets.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: For our experiments on real data, we use publicly available datasets and point to their references [51, 55].

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We provide an implementation of Treeffuser in the supplemental material, which includes an anonymized zip file containing notebooks that document the usage of our method.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.

- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.