
Tensor-Train Density Estimation

Georgii S. Novikov¹

Maxim E. Panov¹

Ivan V. Oseledets¹

¹Skolkovo Institute of Science and Technology, Moscow, Russia

Abstract

Estimation of probability density function from samples is one of the central problems in statistics and machine learning. Modern neural network-based models can learn high dimensional distributions but have problems with hyperparameter selection and are often prone to instabilities during training and inference. We propose a new efficient tensor train-based model for density estimation (TTDE). Such density parametrization allows exact sampling, calculation of cumulative and marginal density functions, and partition function. It also has very intuitive hyperparameters. We develop an efficient non-adversarial training procedure for TTDE based on the Riemannian optimization. Experimental results demonstrate the competitive performance of the proposed method in density estimation and sampling tasks, while TTDE significantly outperforms competitors in training speed.

1 INTRODUCTION

In this paper, we consider a problem of nonparametric density estimation, which is one of the central problems in statistics and machine learning. Recent progress in the development of artificial neural networks has given rise to many new methods of solving this problem, including variational autoencoders (VAE, Kingma and Welling [2014]), generative adversarial networks (GAN, Goodfellow et al. [2014]), autoregressive neural networks [van den Oord et al., 2016], invertible flows [Dinh et al., 2017] among some others. These methods allow us to overcome the curse of dimensionality and make it possible to estimate the density of such high-dimensional and nontrivial data as images and sound. However, all these new approaches lack the simplicity and interpretability of the classical kernel density estimation method [Scott et al., 1977]. On the other hand, kernel den-

sity estimation usually performs poorly even in moderate dimensions [Wang and Scott, 2019].

This paper aims to build a new method of nonparametric density estimation: tensor-train density estimation (TTDE). The idea is to construct a tensor-train approximation to the coefficients' matrix for the expansion of the density function in some basis. We will show that an approximation in this parametric form has several important features that other models do not have (at least not simultaneously): exact sampling, ability to calculate cumulative density function and exact calculation of partition function. Moreover, we propose an efficient training procedure based on Riemannian optimization, which is easy to implement and avoids the problems of instability typical for the methods based on adversarial training.

Contributions of this work. Although extremely powerful and effective, modern neural network-based models have their drawbacks. Some of them do not have tractable log-likelihood at all (like GANs) or have only surrogates for it (lower bound for VAEs, unnormalized log-density for energy-based models). Other methods can not sample from the trained distribution or require the whole additional sampling procedure like MCMC and thus can generate only approximate samples (energy-based models, BNAF [Cao et al., 2019]). Many powerful models require “middle-men” during the training process (discriminator for adversarial models, MCMC sampling for energy-based models), which significantly complicates the development and analysis of such models. These properties for different methods are summarized in Table 1.

Neural network-based methods are famous for their strong dependence on the choice of hyperparameters, architecture and optimization method. We believe that there is a gap between simple, intuitive models for low-dimensional data and powerful, capable of solving most difficult tasks, yet very fragile and hard to theoretically analyze neural network-based methods. In this work, we try to fill this gap.

Table 1: Comparison of the capabilities of different density estimation models. *FFJORD does not use true log-likelihood in the training process and instead uses its unbiased estimate.

Method	Exact Sampling	Tractable LL	No middle-man Training	Computation of CDF
FFJORD	✓	✓*	✓*	✗
Normalizing Flows	✓	✓	✓	✗
GANs	✓	✗	✗	✗
VAEs	✓	✗	✓	✗
Autoregressive	✓	✓	✓	✗
Energy-based	✗	✗	✗	✗
TTDE (ours)	✓	✓	✓	✓

The main contributions of our work are as follows.

- We propose a new generative tensor-based approach *tensor-train density estimation (TTDE)* that allows fast sampling and efficient computation of functionals of probability density function.
- We show that TTDE can be trained using Riemannian optimization targeting a variety of different functionals, including those that are intractable for previously existing models (namely, direct L_2 loss between target probability $p(x)$ and approximation $q_\theta(x)$).
- We illustrate the competitive performance of our approach on a series of examples.

2 WHY TENSOR-TRAIN IS GOOD FOR DENSITY APPROXIMATION?

2.1 PROBLEM STATEMENT

Suppose we are given i.i.d. samples $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}$, $\mathbf{x}^{(i)} \in \mathbb{R}^d, i = 1, \dots, N$ from an unknown probability distribution with a density $p(\mathbf{x})$. We want to find an approximation to this density. It is typically done by using some family of functions:

$$p(\mathbf{x}) \approx q_\theta(\mathbf{x}), \quad (1)$$

where $q_\theta \in \mathcal{Q} = \{q_\theta\}_{\theta \in \Theta}$, $\Theta \subset \mathbb{R}^D$.

To perform approximation (1), some measure of discrepancy between probability densities p and q_θ should be computed (given only samples from p) and then optimized with respect to θ .

In this paper, we propose to use densities represented in the *low-rank tensor-train format* as \mathcal{Q} . This approach has been shown to be successful in [Dolgov et al., 2020] which targeted the problem of the computationally efficient approximation to the given density. This problem is very different from the problem of density estimation from samples which we consider in our work. We aim to fill this gap by develop-

ing a systematic approach for sample-based training of such models.

The general approach in non-parametric statistics is firstly to choose some basis functions $\Phi(\mathbf{x}) = \{f_k(\mathbf{x})\}_{k=1}^K$ (e.g B-splines or Fourier series) and then search for the approximation in the linear space induced by this basis:

$$q_\theta(\mathbf{x}) = \langle \alpha_\theta, \Phi(\mathbf{x}) \rangle = \sum_{k=1}^K \alpha_{\theta,k} f_k(\mathbf{x}),$$

where usually the coefficients vector α_θ simply coincides with the parameter vector θ , i.e. $\alpha_\theta \equiv \theta$.

One of the standard ways to build a multidimensional basis is to take a Cartesian product of several one-dimensional bases, i.e., setting $f_{i_1, \dots, i_d}(\mathbf{x}) = f_{i_1}(x_1) \cdots f_{i_d}(x_d)$ for d -dimensional input \mathbf{x} and some functions $f_1, \dots, f_d: \mathbb{R} \rightarrow \mathbb{R}$. In this case, the coefficients vector α_θ becomes structured as a d -dimensional tensor $\alpha_\theta \in \mathbb{R}^{K^d}$ whose size grows exponentially with the dimension. In this work, we propose to consider only a low-rank subspace of the linear span, i.e., functions $q_\theta(x)$, weight tensor α_θ of which can be represented in low-rank tensor-train format. Such a representation allows achieving linear in the dimension computational and storage costs for operations such as calculation of a function at a point, differentiation and integration.

2.2 PROPOSED REPRESENTATION OF THE DENSITY

Tensor-product basis. Consider a basis set of 1-dimensional functions $\mathcal{B} = \{f_i\}_{i=1}^m, f_i: \mathbb{R} \rightarrow \mathbb{R}$. The construction of the d -dimensional basis set can be done on top of \mathcal{B} as follows:

$$\mathcal{B}^{(d)} = \{f_{i_1, \dots, i_d}\}_{i_1=1, \dots, i_d=1}^{m, \dots, m},$$

where

$$f_{i_1, \dots, i_d}(\mathbf{x}) = f_{i_1}(x_1) \cdots f_{i_d}(x_d): \mathbb{R}^d \rightarrow \mathbb{R}.$$

We are going to approximate the target distribution via the linear function expansion in this basis:

$$\begin{aligned}
q_{\theta}(x) &\in \mathcal{Q} = \text{span } \mathcal{B}^{(d)}, \\
q_{\theta}(x) &= \sum_{i_1=1, \dots, i_d=1}^{m, \dots, m} \alpha_{\theta}^{i_1, \dots, i_d} f_{i_1, \dots, i_d}(x) \\
&= \sum_{i_1=1, \dots, i_d=1}^{m, \dots, m} \alpha_{\theta}^{i_1, \dots, i_d} f_{i_1}(x_1) \cdots f_{i_d}(x_d) \\
&= \sum_{i_1=1, \dots, i_d=1}^{m, \dots, m} \alpha_{\theta}^{i_1, \dots, i_d} \Phi_{i_1, \dots, i_d}(x) = \langle \alpha_{\theta}, \Phi(x) \rangle,
\end{aligned} \tag{2}$$

where the tensor $\Phi(x)$ is a *rank-1 feature map* defined by

$$\begin{aligned}
\Phi(x) &= \mathbf{f}(x_1) \otimes \cdots \otimes \mathbf{f}(x_d), \\
\mathbf{f}(x) &= (f_1(x), \dots, f_m(x)).
\end{aligned}$$

Such a feature map was previously used for the classification problems with tensor-based models in [Cohen et al., 2016, Khurlov et al., 2018, Stoudenmire and Schwab, 2016].

Tensor-train format. In the form above, α_{θ} is a very large tensor of size m^d . To be able to store and interact with this tensor in high dimensions, we will work with the tensors α_{θ} , which can be represented in the tensor-train (TT) format:

$$\alpha_{\theta}^{i_1 \dots i_d} = G_1[\cdot, i_1, \cdot] G_2[\cdot, i_2, \cdot] \cdots G_d[\cdot, i_d, \cdot]. \tag{3}$$

Here G_i are the so-called cores of the tensor-train decomposition, which are 3-dimensional tensors of size $[r_{i-1} \times m \times r_i]$ (with condition $r_0 = r_d = 1$). Here $G_i[\cdot, a, \cdot]$ represents a matrix which is the a -th slice of the core G_i along the second axis. The vector $\mathbf{r} = (r_1, \dots, r_{d-1})$ is called the vector of TT-ranks. Further in the text, unless otherwise stated, we will use a single natural number r to refer to the TT-decomposition with rank $\mathbf{r} = (r, \dots, r)$.

Computation in the TT-format. When all the coefficients are known, the inner product of two tensors in the TT-format can be computed efficiently. This procedure is summarized in Algorithm 1. It requires $\mathcal{O}(dmr_1^2 r_2)$ operations to compute, where r_1 is the maximum of two ranks of given tensors, r_2 is the minimum of two ranks, d is the number of dimensions of the tensors, and m is the maximum size of the dimensions. For example, the evaluation of the function q_{θ} at a point x according to (2) is an inner-product of the weight tensor with a rank-1 tensor and thus requires only $\mathcal{O}(dmr^2)$ time. In a similar way, we can marginalize out some dimensions and therefore compute marginal densities of q_{θ} :

$$q_{\theta}(x_1, \dots, x_{k-1}) = \left\langle \alpha_{\theta}, \Phi(x_1, \dots, x_{k-1}) \otimes \int \Phi(x_k, \dots, x_d) \right\rangle$$

Algorithm 1: Multiplication of two tensors represented in tensor-train format. On each step, we store the contraction of the two prefixes of the lists of cores. Each such contraction can be updated from the previous step in $\mathcal{O}(\max(r_1^2, r_2^2) \min(r_1, r_2))$ time, which gives $\mathcal{O}(d \max(r_1^2, r_2^2) \min(r_1, r_2))$ complexity of the full product.

Result: Inner product of tensors T_1 and T_2 represented in TT-format with cores $\{G_i^{(1)}\}_{i=1}^d$ and $\{G_i^{(2)}\}_{i=1}^d$ and ranks r_1 and r_2 , respectively

Initialize res with $[1 \times 1]$ identity matrix ;

for $p \leftarrow 1$ **to** d **do**

$res \leftarrow \text{einsum}('ix, inj, xny \rightarrow jy', res, G_p^{(1)}, G_p^{(2)})$;

end

return res ;

where

$$\Phi(x_1, \dots, x_{k-1}) = \mathbf{f}(x_1) \otimes \cdots \otimes \mathbf{f}(x_{k-1}),$$

$$\int \Phi(x_k, \dots, x_d) = \left(\int \mathbf{f}(x_k) dx_k \right) \otimes \cdots \otimes \left(\int \mathbf{f}(x_d) dx_d \right)$$

or even calculate cumulative density function along some dimension k :

$$q_{\theta}(x_1, \dots, x_{k-1}, x_k < A) =$$

$$\left\langle \alpha_{\theta}, \Phi(x_1, \dots, x_{k-1}) \otimes \int_{-\infty}^A f(x_k) dx_k \otimes \int \Phi(x_{k+1}, \dots, x_d) \right\rangle.$$

Squared TTDE. The main drawback of proposed model is that it's not guaranteed to be always non-negative. If the expressivity of the model is sufficient, we expect the resulting approximation to be sufficiently similar to the real distribution, and hence the negative regions can be neglected. However, it is too optimistic to expect such behavior for complex high-dimensional distributions. If we want to apply the model to medium dimensions, we have to overcome this problem. One solution is to require that all the kernels in the tensor train (3) are non-negative. Then, if all basis functions are also always non-negative (which can be assumed by construction), then the resulting function q_{θ} will be non-negative everywhere. Notice, that this approach will require additional changes to the optimization process. We propose another modification of our model – *squared TTDE*. Instead of approximation (3) we suggest to use squared version of it:

$$q_{\theta}(x) = \left\langle \alpha_{\theta}, \Phi(x) \right\rangle^2,$$

which automatically implies non-negativity. In this form, some operations will become more computationally expensive.

sive (e.g., calculating marginals), but still manageable, while the model can be trained using the classical method of likelihood maximization. Experiments with the squared TTDE are presented in Section 5.4 and summarized in Table 3.

2.3 SAMPLING

Direct application of the trained density model is sampling from that model. GANs, for example, can not do anything but sampling (and require one forward pass through the network to do it). Models based on normalizing flows can both infer density and sample (and in general require one forward pass as well). Energy-based models can not generate exact samples from their learned densities and moreover, require an additional iterative procedure to get approximate samples (like Markov Chain Monte Carlo sampling).

Probability density function represented in the tensor-train format allows fast, exact sampling in the autoregressive fashion: as we can calculate marginals of $q_\theta(x)$, we can take sample seeds $\mathbf{u} \sim \mathcal{U}([0; 1]^d)$ and then sample coordinates of \mathbf{x} one by one such that $q_\theta(\xi_1 < x_1) = u_1$, $q_\theta(\xi_2 < x_2 \mid \xi_1 = x_1) = u_2$ and so on. On a step k to sample x_k , we assume that we already know all $x_i, i < k$ and thus should find such a number A , that

$$q_\theta(\xi_k < A \mid \{\xi_i = x_i\}_{i < k}) = \frac{q_\theta(\xi_k < A, \{\xi_i = x_i\}_{i < k})}{q_\theta(\{\xi_i = x_i\}_{i < k})} = u_k,$$

which is a simple 1-dimensional search on a monotonically increasing function (up to an approximation error) and thus can be performed with any appropriate algorithm (e.g. binary search).

TT-format allows an efficient implementation of the described algorithm. Notice that the cumulative density function value can be computed in a cycle over dimensions, where each iteration k of which can be decomposed in four steps:

1. Contraction of $\Phi(x_1, \dots, x_{k-1})$ with cores G_1, \dots, G_{k-1} . Let us call it Q_k^{left} , which is a vector of size r_{k-1} .
2. Contraction of $\int \Phi(x_{k+1}, \dots, x_d) dx_{k+1} \dots dx_d$ with cores G_{k+1}, \dots, G_d . Lets call it Q_k^{right} , which is a vector of size r_k .
3. Multiplication of core G_k with Q_k^{left} and Q_k^{right} along it's left and right dimensions, which results in a vector Q_k^{inner} of size m , which is the size of the 1-dimensional basis.
4. Dot product of Q_k^{inner} with the vector $\int_{-inf}^{x_k} f(x_k) dx_k$.

In the considered algorithm, Q_k^{right} can be precomputed beforehand. Q_k^{left} depends only on x_i for $i < k$, and con-

Algorithm 2: Algorithm to retrieve an exact sample \mathbf{x} from the density $q_\theta(\mathbf{x})$ represented in TT-format.

Result: Sample $\mathbf{x} \sim q_\theta$

Sample $\mathbf{u} \sim \mathcal{U}([0; 1]^d)$;

Precompute Q_k^{right} for all k ;

Initialize Q_1^{left} with 1 ;

for $k \leftarrow 1$ **to** d **do**

Precompute vector Q_k^{inner} ;

Find such x_k , that

$$\left(Q_k^{inner}, \int_{-inf}^{x_k} \mathbf{f}(x_k) dx_k \right) = u_k ;$$

Update Q_k^{left} ;

return \mathbf{x}

sequently, Q_{k+1}^{left} can be derived from Q_k^{left} by one matrix-vector product. Thus, steps 1, 2 and 3 require only $\mathcal{O}(dmr^2)$ operations in total. That means that one calculation of the cumulative density function during the 1-dimensional search boils down to a cheap calculation of $\int_{-inf}^{x_k} \mathbf{f}(x_k) dx_k$ and one vector-vector product of size m . In total, this algorithm requires $\mathcal{O}(dmr^2 + dmL)$ operations, where L is a number of iterations that a 1-dimensional search would do (for binary search L would be proportional to the log of the desired result precision). The resulting sampling procedure is summarized in Algorithm 2.

3 LEARNING VIA RIEMANNIAN OPTIMIZATION

Loss function. In practice, we do not know the true probability function p and have access only to the dataset $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^n$ of i.i.d. samples \mathbf{x}_i from density p . Thus, we need to optimize some loss function in order to get an approximation q_θ of p . In the previous section we showed that we can explicitly calculate the partition function for the proposed tensor train-based model. Similarly, it is possible to calculate L_2 norm of the function represented in the tensor-train format in just a $\mathcal{O}(d(r^3m + r^2m^2))$ time. That allows us to use an interesting loss, unusual in the density estimation context, computed by the direct calculation of L_2 distance between target distribution p and approximation q_θ :

$$\begin{aligned} \mathcal{L}(p, q_\theta) &= \int (p(\mathbf{x}) - q_\theta(\mathbf{x}))^2 dx \\ &= \int q_\theta(\mathbf{x})^2 dx - 2 \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} q_\theta(\mathbf{x}) + const. \end{aligned} \quad (4)$$

As only samples are available from density p , the expectation in (4) can be approximate with the Monte-Carlo estimate based on the samples from the dataset. If the expressive power of q_θ is large enough, then minimum will be achieved near p and thus will produce a good approximation to the

true density function.

Computation of loss function and its derivatives. The second term in (4) is just an evaluation of the function at a given point. It was discussed in details in Section 2.2 and can be calculated in $\mathcal{O}(bdr^2m)$ time, where b is a batch size used in the stochastic optimization method. The first term in (4) is a quadratic function w.r.t. tensor α_θ and thus can be expressed in form $\langle \alpha_\theta, D\alpha_\theta \rangle$ with an appropriate choice of the linear operator D :

$$\begin{aligned} \int \langle \alpha_\theta, \Phi(\mathbf{x}) \rangle^2 d\mathbf{x} &= \int \langle \alpha_\theta, (\Phi(\mathbf{x}) \otimes \Phi(\mathbf{x})) \alpha_\theta \rangle d\mathbf{x} \\ &= \left\langle \alpha_\theta, \left(\int \Phi(\mathbf{x}) \otimes \Phi(\mathbf{x}) d\mathbf{x} \right) \alpha_\theta \right\rangle, \end{aligned}$$

where

$$\begin{aligned} &\left[\int \Phi(\mathbf{x}) \otimes \Phi(\mathbf{x}) d\mathbf{x} \right]_{i_1, j_1, \dots, i_d, j_d} \\ &= \int \Phi(\mathbf{x})_{i_1, \dots, i_d} \otimes \Phi(\mathbf{x})_{j_1, \dots, j_d} d\mathbf{x} \\ &= \int f_{i_1}(x_1) f_{j_1}(x_1) dx_1 \cdots \int f_{i_d}(x_d) f_{j_d}(x_d) dx_d \\ &= D_{i_1, j_1} \cdots D_{i_d, j_d} \end{aligned}$$

and

$$D_{i,j} = \int f_i(x) f_j(x) dx.$$

Thus, tensor $\left[\int \Phi(\mathbf{x}) \otimes \Phi(\mathbf{x}) d\mathbf{x} \right]$ is rank-1 tensor being an outer product of d matrices $D_{i,j}$. Thus, an application of it to the tensor α_θ boils down to the multiplication of each core of α_θ along the middle axis with the matrix D followed by the inner product of two rank- r tensors. It results in total complexity $\mathcal{O}(dmr^2(m+r))$. Let us note that an additional multiplication by $(m+r)$ does not significantly increase the computational cost of the optimization process as this term does not depend on the batch size b .

Riemannian optimization and optimal step. In principle, the standard stochastic gradient descent methods or their variations can be used to train the proposed model. However, the representation of the model in tensor-train format and the fact, that the given loss function (4) is quadratic with respect to the tensor α allow for usage of more productive optimization methods. In this work, we suggest using Riemannian optimization, which is a promising tool for learning tensor-based models [Rakhuba et al., 2019, Steinlechner, 2016].

Riemannian optimization is a procedure to minimize some function g defined on \mathcal{X} over some smooth manifold $\mathcal{M} \subset$

\mathcal{X} :

$$\min_{\mathbf{X} \in \mathcal{M}} g(\mathbf{X}).$$

The usual Riemannian optimization workflow consists of several steps:

1. Construction of *tangential plane* $\mathcal{T}_{\mathbf{x}}(\mathcal{M})$ to manifold \mathcal{M} at point \mathbf{x} . For tensor-train format, it has an efficiently computable expression, see [Rakhuba et al., 2019].
2. Projection $\mathbb{P}_{\mathcal{T}_{\mathbf{x}}(\mathcal{M})} \nabla g(\mathbf{X})$ of the true gradient $\nabla g(\mathbf{X})$ onto the tangent plane $\mathcal{T}_{\mathbf{x}}(\mathcal{M})$, which can be done efficiently using automatic differentiation.
3. Gradient step in the tangent plane: $\mathbf{X}^{next} = \mathbf{X} + \alpha \mathbb{P}_{\mathcal{T}_{\mathbf{x}}(\mathcal{M})} \nabla g(\mathbf{X})$, where α is a learning rate.
4. Retraction of a point \mathbf{X}^{next} back onto \mathcal{M} , which again can be efficiently approximated for tensor-train format.

By the construction, the tangent plane $\mathcal{T}_{\mathbf{x}}(\mathcal{M})$ is a linear space, and due to the fact that our loss function (4) is a quadratic function, we can find optimal α on each step as a minimal point of a parabola. Note that it is not true for the classical gradient descent in the space of TT-cores, as there will be a complex high degree polynomial dependence.

More details on how to construct $\mathcal{T}_{\mathbf{x}}(\mathcal{M})$, how to represent it in the tensor-train format of doubled rank $2r$ and how to project the true gradient onto tangent space can be found in Supplementary Material A.3 and in [Rakhuba et al., 2019].

Initialization. It is important to have good initialization for the gradient optimization methods in general as well as for the proposed Riemannian optimization approach. Rather efficient but straightforward initialization can be performed under the assumption of coordinate independence. Consider the case where

$$p(\mathbf{x}) = \prod_{i=1}^d p_i(x_i)$$

for some set of one-dimensional probability density functions $\{p_i\}_{i=1}^d$. We can solve approximation problems

$$p_i(x_i) \approx \langle \alpha_i, f(x_i) \rangle, \quad i = 1, \dots, d \quad (5)$$

independently in such a case and then consider $\alpha_\theta = \alpha_1 \otimes \cdots \otimes \alpha_d$ as a rank-1 tensor-train initialization. Each approximation (5) can be computed as the solution of a simple linear regression problem for the loss (4).

4 RELATED WORK

One of the most famous non-parametric density estimation algorithms is the histogram approach [Scott, 1979], which

works only in very low dimensional settings (1-2 dimensions). Another famous method is the celebrated kernel density estimation approach [Scott et al., 1977] that is again known to perform poorly in high dimensions. An important quality of these classical methods for approximating the distribution of low-dimensional data is their simplicity and intuitive behavior. However, many modern methods discussed below greatly improve the quality of density estimation which comes at the cost of much more sophisticated estimation procedures.

The recent development of artificial neural networks gave birth to several new families of non-parametric density estimation. Generative-adversarial networks (GANs, Goodfellow et al. [2014]) are methods of building the neural network capable of generating synthetic data close to the observed data. Although astonishing performance in real-life problems and the ability to learn and generalize extremely high-dimensional and complex data (images, videos and sound), this methods do not produce tractable (or even intractable) density functions and thus the applicability of these methods in statistical context is limited. Another similar approach is variational autoencoder (VAE, Kingma and Welling [2014]). Unlike GANs, VAEs minimize the variational lower bound of the likelihood and thus can be used to approximate the unnormalized density functions, although the partition function is still intractable.

There is also a great variety of methods, based on neural networks, that directly learn the density function: energy-based models [LeCun et al., 2006], autoregressive density estimators Ryder et al. [2018] and normalizing flows Kobyzev et al. [2020], among some others. Energy-based models learn the unnormalized density functions by maximizing the log-likelihood of the data and approximate the partition function by MCMC sampling. Thus, only approximate sampling from these models is available (via MCMC). Autoregressive models factor the probability distribution $p(x) = \prod_{k=1}^d p(x_k | x_{s < k})$, parameterizing each factor with a neural network. Methods based on normalizing flows build smooth bijection of the target space with the latent space of the same size: $z = f_{\theta}(x) \Rightarrow p(x) = p(z) \left| \frac{\partial f}{\partial x} \right|$. By setting the simple distribution of the latent variables (usually standard Gaussian) and assuming that the log-determinant of the function f_{θ} can be efficiently calculated (guaranteed by choosing the appropriate neural network structure), the likelihood of the observed data can be directly maximized. This method was successfully applied to such complex tasks as face and speech generation [Kingma and Dhariwal, 2018, Kim et al., 2020]. The recent development of NF models allows the use of more and more complex and less constrained models (see, for example, FFJORD [Grathwohl et al., 2019]), but at the same time, some of them lose the ability to sample from the trained model (BNAF [Cao et al., 2019]).

Another family of models that offers tractable query class is Probabilistic Circuits [Choi et al.] – acyclic directional computational graphs that represent complex distributions as mixtures (sum nodes in the graph) and factorizations (product nodes in the graph) with simple tractable distributions, usually one-dimensional, in the leaves of this graphs (input nodes). To some extent, TTDE can be seen as a Probabilistic Circuit with basis functions being input nodes, and the graph structure implicitly defined in agreed with tensor-train format (3), which would give us, although predefined, *exponentially large* model structure. Also, TTDE admits negative weights, which potentially beneficially affects expressiveness [Dennis, 2016].

5 EXPERIMENTS

We evaluate the performance of the presented TTDE method on several model and real-world datasets. All the code to reproduce the results of experiments can be found via <https://github.com/stat-ml/TTDE>. In all the experiments basis function set consists of B-splines of degree 2 with knots uniformly distributed over the support of the considered distributions. The support is known precisely for the simulated examples as we know exactly the target distribution, and corresponding lower and upper bounds are extracted from all given samples for the unknown distributions of the real-world datasets.

5.1 TOY EXAMPLES

We start the presentation of our results with the several classical 2-dimensional examples, which are presented on the Figure 1. We clearly observe that TTDE doesn't have any problems with complex shape of the distribution, its multimodality or discontinuity. Moreover, the visual comparison shows clear superiority of TTDE over state-of-the-art normalizing flow approach FFJORD [Grathwohl et al., 2019].

5.2 MODEL DATASETS AND HYPERPARAMETER SELECTION

An important feature of our model is great interpretability of the model hyperparameters: basis size and tensor-train rank. The basis size corresponds to the resolution of the approximation. It acts similarly to the number of bins in the very large multidimensional histogram. Rank of the tensor-train decomposition corresponds to the expressive power of the model, i.e. how complex distributions can be built for the given basis size. The dependence of the trained density on both hyperparameters is shown on Figure 2 for the celebrated two moons dataset. We clearly observe the behavior discussed above. Interestingly, if the rank of tensor-train decomposition is not large enough, the method tries to somehow cope with it, adding symmetric artifacts to the

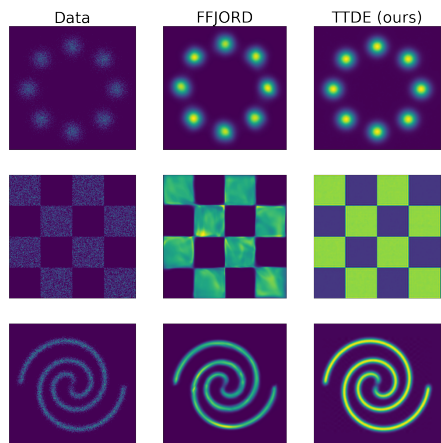


Figure 1: Comparison of TTDE and FFJORD models on 2-dimensional toy distributions.

	Random init.	Rank-1 init.
Adam	5	11
Riemannian	12	32

Table 2: Experiment with mixture of 7 Gaussians in 3D with additional dimensions containing only noise. We report the maximum dimensionality for which approximation of the density converges to the true one for different initialization settings and optimization methods used.

distribution.

We additionally explore TTDE properties dependence on the rank of the model for the 8-dimensional mixture of 128 Gaussians located in the random corners of 8-dimensional unit cube. We present the dependence of the cross-entropy on the rank of the tensor-train decomposition on Figure 3. We observe an expected behavior: the higher is the rank, the higher is the cross entropy with the true distribution. Importantly, already the approximation of the rank 16 is enough to almost perfectly match this very complex distribution.

5.3 IMPORTANCE OF INITIALIZATION.

To show the importance of the proper initialization and optimization method, we tested our method on the mixture of seven identical n -dimensional Gaussian mixtures with identity covariance matrix and located in the 7 corners (all except one) of the 3-dimensional cube in the first 3 dimensions of the space (see the illustration in Supplementary Material). This is rather simple distribution having rank 2. We trained four different models: with rank-1 initialization (see Section 3) or with random initialization, and with proposed Riemannian optimizer (see Section 3) or with the standard Adam optimizer. We set the rank of the approximation equal to 4 (slightly larger, than the true rank of the target distribution).

For each combination of initialization and optimization method we report the largest dimensionality of the data for which the corresponding method successfully converges to correct solution. The resulting numbers are shown in Table 2. We see that for random initialization Riemannian optimization allows to achieve correct results in much higher dimension than Adam (12 vs 5). The usage of the proposed 1d-initialization procedure allows to significantly boost the results for both optimizers. The clear winner is properly initialized model with Riemannian optimization which is capable to learn 32– dimensional distribution. Thus if we want to apply this algorithm in high dimensional settings, the proposed initialization and the right choice of optimization technique are vital.

5.4 REAL-WORLD DATA

Methods and data. To show the applicability of our method to the real-world tasks, we compare computational performance and quality of approximation of TTDE and squared-TTDE with several methods from the Normalizing Flows family as they have similar capabilities (see Table 1), namely Glow [Kingma and Dhariwal, 2018], Real-NVP [Dinh et al., 2017], MAF [Papamakarios et al., 2017a] and FFJORD [Grathwohl et al., 2019]. We perform comparison on five tabular datasets from UCI dataset collection preprocessed as in Papamakarios et al. [2017b].

Quality measure. Due to the fact that the function represented in the tensor-train format does not have to be a valid probability density function i.e. potentially there could be areas of small negative values due to approximation error, direct comparison in terms of model likelihood is not available. On the other hand, comparison in terms of the loss (4) is not a good choice as well. Firstly, it is problematic to calculate it for the normalizing flow models in 6-dimensional space. Secondly, it is not fair to compare the model that directly optimizes this loss (our model) with the models, that are trained using completely different discrepancy measures. Because of that, we decided to compare models based on the quality of generated samples. Namely, we decided to measure the *sliced total variation* between the samples from the validation set and the samples acquired from the model.

Total variation is a classical measure of discrepancy between two distributions

$$TV(p \parallel q) = \int |p(x) - q(x)| dx.$$

In spite of the simplicity of its formulation, the computation of the integral above in high dimension is hard and potentially non-accurate. That’s why we introduce the sliced version of the total variation by employing a simple idea how to apply it to the multidimensional case when the true distribution $p(x)$ is not available. Instead of integrating over

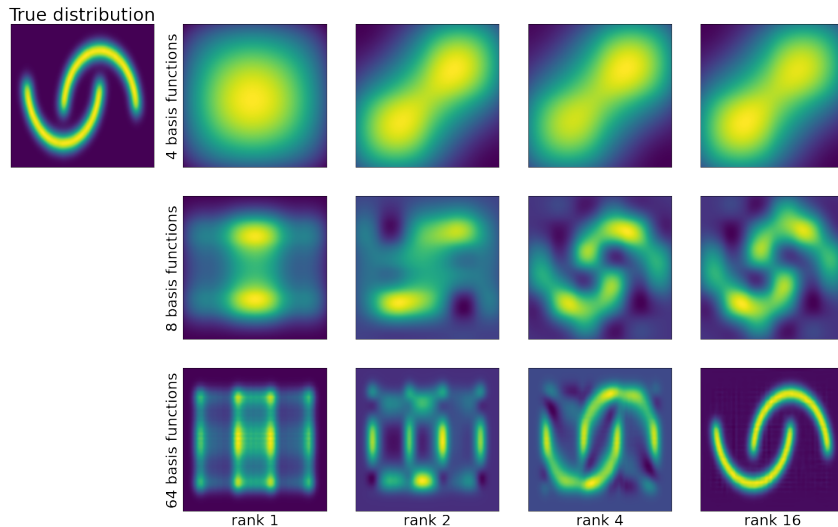


Figure 2: Approximations of “two moons” distribution by TTDE for different basis function set sizes and TT-ranks.

Table 3: Average log-likelihood for several neural-network-based models on tabular UCI datasets. Gaussian fitted to the training data is reported as a baseline.

*On Hepmass and Miniboone datasets, which has the lowest number of training examples (300k and 30k respectively), we observe heavy overfitting. Lack of regularizations for the new model leads to poor results. Thus, it is an important direction for further development of the TTDE.

	POWER	GAS	HEPMASS	MINIBOONE	BSDS300
Dataset dimensionality	6	8	21	43	64
Gaussians	-7.74	-3.58	-27.93	-37.24	96.67
MADE	-3.08	3.56	-20.98	-15.59	148.85
Real NVP	0.17	8.33	-18.71	-13.84	153.28
Glow	0.17	8.15	-18.92	-11.35	155.07
FFJORD	0.46	8.59	-14.92	-10.43	157.40
Squared TTDE (ours)	0.46	8.93	-21.34*	-28.77*	143.30

the whole space, we can average total variations for many random 1-dimensional projections:

$$STV(p \parallel q) = \mathbb{E}_P \int |P[p](x) - P[q](x)| dx,$$

where P is a random projector on one dimensional space. In our experiments we firstly generate samples from all testing models. Then several times we generate projection of validation set and all generated sets on a random 1-dimensional plane. Then we calculate 1-dimensional approximations of distributions using kernel density estimation (which almost exactly replicates true underlying distribution due to massive amounts of generated points) and then we calculated TV of this 1-dimensional functions. This procedure is summarized in Algorithm 3.

Results. We show in our experiments that the proposed model significantly outperforms (by an order of magnitude)

all presented neural-network based models in terms of the convergence rate to the optimal value, see Figure 4. The same applies to the speed of sampling, see Figure 5, where for the batch size of 2^{20} we outperform two most powerful baseline models: FFJORD and MAF (speedup of 2.8 and 2.5 times respectively) and slightly outperform GLOW and Real NVP (1.4 and 1.2 times speedup respectively).

In table 3 we report log-likelihoods achieved with our squared-TTDE model trained with NLL loss and compare them with several normalizing flow models. Our model manages to outperform presented competitors on two low-dimensional datasets (POWER and GAS). On the 64-dimensional BSDS300 dataset, our model performs worse than powerful network-based models, although by not much, while still providing tractable query class. On Hepmass and Miniboone datasets, which have the lowest number of training examples (300k and 30k respectively), we observe heavy overfitting. We think that development of regularization tech-

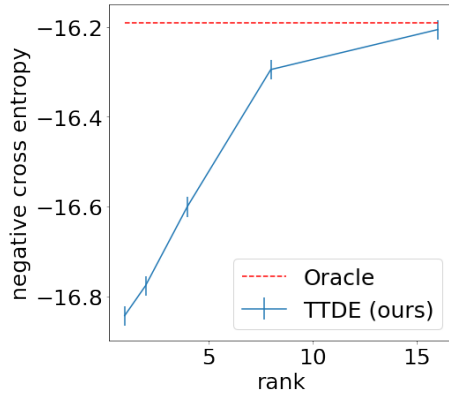


Figure 3: Dependence of the approximation quality of the mixture of 128 Gaussians in 8 dimensional space on the rank of the TT decomposition.

Algorithm 3: Calculation of sliced total variation between two distributions based on samples from them.

Result: Sliced total variation for two sample sets

$$X^{(1)} = \{x_i^{(1)} \in \mathbb{R}^d\}_{i=1}^N \text{ and}$$

$$X^{(2)} = \{x_i^{(2)} \in \mathbb{R}^d\}_{i=1}^N$$

Choose random 1 - dimensional hyperplane l ;

Project $X^{(1)}$ and $X^{(2)}$ onto l ;

Build approximations p_1 and p_2 from projected samples;

Approximate $STV = \int_{\mathbb{R}} |p_1(y) - p_2(y)| dy$ using any 1-dimensional numerical integration method ;

return STV ;

niques for the TTDE will allow to obtain better results. All the details of experiments are specified in Supplementary Materials A.1.

6 CONCLUSION

This work shows that approximation based on the tensor-train decomposition is a promising method of density estimation. It offers such a set of different features and possibilities (tractable partition function, exact sampling, exact marginals, and cumulative densities), which were not previously accessible all at the same time for other methods. This method’s ability to work in medium dimensionality is very promising and paves the way to accurate density estimation in high dimensions.

Acknowledgements

The research was supported by the Russian Science Foundation grant 20-71-10135.

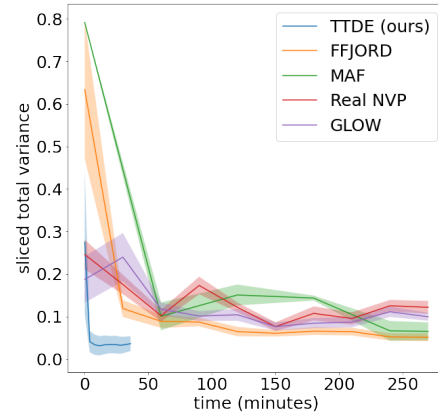


Figure 4: Dependence of the sliced total variation w.r.t. the training time for models trained on 6-dimensional UCI POWER dataset.

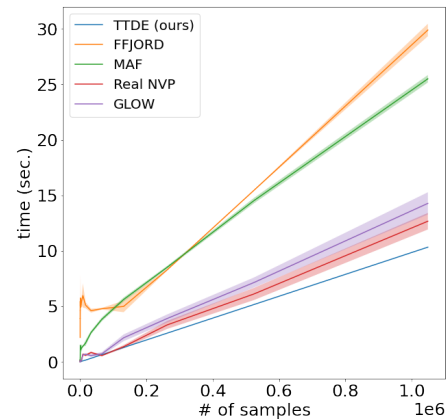


Figure 5: Dependence of the sampling time w.r.t. the number of samples to be generated for 6-dimensional space for models trained on UCI POWER dataset. Our model outperforms its competitors and shows 2.6, 2.5, 1.4 and 1.2 times speedups compared to FFJORD, MAF, GLOW and Real NVP respectively.

References

- Nicola De Cao, Wilker Aziz, and Ivan Titov. Block neural autoregressive flow. In Amir Globerson and Ricardo Silva, editors, *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, July 22-25, 2019*, volume 115 of *Proceedings of Machine Learning Research*, pages 1263–1273. AUAI Press, 2019. URL <http://proceedings.mlr.press/v115/de-cao20a.html>.
- YooJung Choi, Antonio Vergari, and Guy Van den Broeck. Probabilistic circuits: A unifying framework for tractable probabilistic models.
- Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analy-

- sis. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *Proceedings of the 29th Conference on Learning Theory, COLT 2016, New York, USA, June 23-26, 2016*, volume 49 of *JMLR Workshop and Conference Proceedings*, pages 698–728. JMLR.org, 2016. URL <http://proceedings.mlr.press/v49/cohen16.html>.
- Aaron W Dennis. Algorithms for learning the structure of monotone and nonmonotone sum-product networks. 2016.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=HkpbH91lx>.
- Sergey Dolgov, Karim Anaya-Izquierdo, Colin Fox, and Robert Scheichl. Approximation and sampling of multivariate probability distributions in the tensor train decomposition. *Stat. Comput.*, 30(3):603–625, 2020. doi: 10.1007/s11222-019-09910-z. URL <https://doi.org/10.1007/s11222-019-09910-z>.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>.
- Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. FFJORD: free-form continuous dynamics for scalable reversible generative models. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rJxgknCcK7>.
- Valentin Khruikov, Alexander Novikov, and Ivan V. Oseledets. Expressive power of recurrent neural networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=S1WRibb0Z>.
- Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sungroh Yoon. Glow-tts: A generative flow for text-to-speech via monotonic alignment search. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/5c3b99e8f92532e5ad1556e53ceea00c-Abstract.html>.
- Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 10236–10245, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/d139db6a236200b21cc7f752979132d0-Abstract.html>.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing Flows: An Introduction and Review of Current Methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, May 2020. ISSN 0162-8828. doi: 10.1109/tpami.2020.2992934.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Marc’Aurelio Ranzato, and Fu-Jie Huang. A tutorial on energy-based learning. In G. Bakir, T. Hofman, B. Schölkopf, A. Smola, and B. Taskar, editors, *Predicting Structured Data*. MIT Press, 2006.
- George Papamakarios, Iain Murray, and Theo Pavlakou. Masked autoregressive flow for density estimation. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2338–2347, 2017a. URL <https://proceedings.neurips.cc/paper/2017/hash/6c1da886822c67822bcf3679d04369fa-Abstract.html>.
- George Papamakarios, Iain Murray, and Theo Pavlakou. Masked autoregressive flow for density estimation. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural*

- Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2338–2347, 2017b. URL <https://proceedings.neurips.cc/paper/2017/hash/6c1da886822c67822bcf3679d04369fa-Abstract.html>.
- Maxim V. Rakhuba, Alexander Novikov, and Ivan V. Osleledets. Low-rank riemannian eigensolver for high-dimensional hamiltonians. *J. Comput. Phys.*, 396:718–737, 2019. doi: 10.1016/j.jcp.2019.07.003. URL <https://doi.org/10.1016/j.jcp.2019.07.003>.
- Tom Ryder, Andrew Golightly, A. Stephen McGough, and Dennis Prangle. Black-box autoregressive density estimation for state-space models. *CoRR*, abs/1811.08337, 2018. URL <http://arxiv.org/abs/1811.08337>.
- David W. Scott. On optimal and data-based histograms. *Biometrika*, 66(3):605–610, dec 1979. ISSN 0006-3444. doi: 10.1093/biomet/66.3.605. URL <https://academic.oup.com/biomet/article-lookup/doi/10.1093/biomet/66.3.605>.
- David W. Scott, Richard A. Tapia, and James R. Thompson. Kernel density estimation revisited. *Nonlinear Analysis*, 1(4):339–372, jan 1977. ISSN 0362546X. doi: 10.1016/S0362-546X(97)90003-1.
- Michael Steinlechner. Riemannian optimization for high-dimensional tensor completion. *SIAM J. Sci. Comput.*, 38(5), 2016. doi: 10.1137/15M1010506. URL <https://doi.org/10.1137/15M1010506>.
- Edwin Miles Stoudenmire and David J. Schwab. Supervised learning with tensor networks. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 4799–4807, 2016. URL <https://proceedings.neurips.cc/paper/2016/hash/5314b9674c86e3f9d1ba25ef9bb32895-Abstract.html>.
- Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1747–1756. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/oord16.html>.
- Zhipeng Wang and David W. Scott. Nonparametric density estimation for high-dimensional data - algorithms and applications. *CoRR*, abs/1904.00176, 2019. URL <http://arxiv.org/abs/1904.00176>.