
Deep Unsupervised Learning using Nonequilibrium Thermodynamics

Jascha Sohl-Dickstein

Stanford University

JASCHA@STANFORD.EDU

Eric A. Weiss

University of California, Berkeley

EWEISS@BERKELEY.EDU

Niru Maheswaranathan

Stanford University

NIRUM@STANFORD.EDU

Surya Ganguli

Stanford University

SGANGULI@STANFORD.EDU

Abstract

A central problem in machine learning involves modeling complex data-sets using highly flexible families of probability distributions in which learning, sampling, inference, and evaluation are still analytically or computationally tractable. Here, we develop an approach that simultaneously achieves both flexibility and tractability. The essential idea, inspired by non-equilibrium statistical physics, is to systematically and slowly destroy structure in a data distribution through an iterative forward diffusion process. We then learn a reverse diffusion process that restores structure in data, yielding a highly flexible and tractable generative model of the data. This approach allows us to rapidly learn, sample from, and evaluate probabilities in deep generative models with thousands of layers or time steps.

1. Introduction

Historically, probabilistic models trade-off between two conflicting objectives: *tractability* and *flexibility*. Models that are *tractable* can be analytically evaluated and easily fit to data (e.g. a Gaussian or Laplace). However, these models are unable to aptly describe structure in rich datasets. On the other hand, models that are *flexible* can be molded to fit structure in arbitrary data. For example, we can define models in terms of any (non-negative) function $\phi(\mathbf{x})$ yielding the flexible distribution $p(\mathbf{x}) = \frac{\phi(\mathbf{x})}{Z}$, where Z is a normalization constant. However, computing this normalization constant is generally intractable. Evaluating, training,

or drawing samples from flexible models typically requires a very expensive Monte Carlo process.

A variety of analytic approximations exist which ameliorate, but do not remove, this tradeoff – for instance mean field theory and its expansions [36, 37], variational Bayes [17], contrastive divergence [39, 12], minimum probability flow [33, 32], minimum KL contraction [25], proper scoring rules [8, 29], score matching [14], pseudolikelihood [4], loopy belief propagation [26], and many, many more. Non-parametric methods [7] can also be very effective¹.

1.1. Diffusion probabilistic models

We present a way to define probabilistic models that allows:

1. extreme flexibility in model structure,
2. exact sampling,
3. easy multiplication with other distributions, e.g. in order to compute a posterior, and
4. the model log likelihood, and the probability of individual states, to be cheaply evaluated.

Our method uses a Markov chain to gradually convert one distribution into another, an idea used in non-equilibrium statistical physics [15] and sequential Monte Carlo [27]. We build a Markov chain which converts a simple known distribution (e.g. a Gaussian) into a target (data) distribution using a diffusion process. Rather than use this Markov chain to approximately evaluate a model which has been

¹Non-parametric methods can be seen as transitioning smoothly between tractable and flexible models. For instance, a non-parametric Gaussian mixture model will represent a small amount of data using a single Gaussian, but may represent infinite data as a mixture of an infinite number of Gaussians.

otherwise defined, we explicitly define the probabilistic model as the endpoint of the Markov chain. Since each step in the diffusion chain has an analytically evaluable probability, the full chain can also be analytically evaluated.

Learning in this framework involves estimating small perturbations to a diffusion process. Estimating small perturbations is more tractable than explicitly describing the full distribution with a single, non-analytically-normalizable, potential function. Furthermore, since a diffusion process exists for any smooth target distribution, this method can capture data distributions of arbitrary form.

We demonstrate the utility of these *diffusion probabilistic models* by training high log likelihood models for a two-dimensional swiss roll, binary sequence, handwritten digit (MNIST), and several natural image (CIFAR-10, bark, and dead leaves) datasets.

1.2. Relationship to other work

The wake-sleep algorithm [13, 5] introduced the idea of training inference and generative probabilistic models against each other. This approach remained largely unexplored for nearly two decades, with some exceptions [31, 18]. There has been a recent explosion of work developing this idea. In [19, 10, 30, 28], variational learning and inference algorithms were developed which allow a flexible generative model and posterior distribution over latent variables to be directly trained against each other. The variational bound in these papers is similar to the one used in our training objective and in the earlier work of [31].

However, our motivation and model form are both quite different, and the present work retains the following differences and advantages relative to these techniques:

1. We derive a framework using ideas from physics, quasi-static processes, and annealed importance sampling rather than as stemming from variational approximation methods.
2. We address the difficulty that training the inference model often proves particularly challenging in variational inference methods. We restrict the forward (inference) process to a simple functional form, in such a way that the reverse (generative) process will have the same functional form.
3. We train models with thousands of layers (or time steps), rather than only a handful of layers.
4. We show how to easily multiply the learned distribution with another probability distribution (eg with a conditional distribution in order to compute a posterior)
5. We provide upper and lower bounds on the entropy production in each layer (or time step)
6. We provide state of the art unsupervised learning re-

sults on an extremely complex dataset, natural images.

Other related methods for training probabilistic models include generative stochastic networks [2, 40] which directly train a Markov kernel to match its equilibrium distribution to the data distribution, adversarial networks [9] which train a generative model against a classifier which attempts to distinguish generated samples from true data, and mixtures of conditional Gaussian scale mixtures (MCGSMs) [38] which describe a dataset using Gaussian scale mixtures, with parameters which depend on a sequence of causal neighborhoods. We compare experimentally against adversarial networks and MCGSMs.

Related ideas from physics include the Jarzynski equality [15] (known in machine learning as annealed importance sampling [27]), which uses a Markov chain which converts one distribution into another to compute the ratio of normalizing constants between the two distributions; Langevin dynamics [21] which show how to define a Gaussian diffusion process which has any target distribution as its equilibrium; and the Kolmogorov forward and backward equations [6] which show that forward and reverse diffusion processes can be described using the same functional form.

2. Algorithm

Our goal is to define a forward diffusion process which converts any complex data distribution into a simple, tractable, distribution, and then to use the reversal of this diffusion process to define our generative model distribution. We first describe the forward diffusion process, and the reversal of this diffusion process. We then show how these diffusion processes can be trained and used to evaluate probabilities. We also derive entropy bounds for the reverse processes, and show how the learned distributions can be multiplied by any second distribution (e.g. as would be done to compute a posterior when inpainting or denoising an image).

2.1. Forward Trajectory

We label the data distribution $q(\mathbf{x}^{(0)})$. The data distribution is gradually converted into a well behaved (analytically tractable) distribution $\pi(\mathbf{y})$ by repeated application of a Markov diffusion kernel $T_\pi(\mathbf{y}|\mathbf{y}';\beta)$ for $\pi(\mathbf{y})$, where β is the diffusion rate,

$$\pi(\mathbf{y}) = \int d\mathbf{y}' T_\pi(\mathbf{y}|\mathbf{y}';\beta) \pi(\mathbf{y}') \quad (1)$$

$$q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) = T_\pi(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)};\beta_t). \quad (2)$$

The forward trajectory, corresponding to starting at the data distribution and performing T steps of diffusion, is thus

$$q(\mathbf{x}^{(0 \dots T)}) = q(\mathbf{x}^{(0)}) \prod_{t=1}^T q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \quad (3)$$




Figure 1. The proposed modeling framework trained on 2-d swiss roll data. The top row shows time slices from the forward trajectory $q(\mathbf{x}^{(0\dots T)})$. The data distribution (left) undergoes Gaussian diffusion, which gradually transforms it into an identity-covariance Gaussian (right). The bottom row shows the corresponding time slices from the trained reverse trajectory $p(\mathbf{x}^{(0\dots T)})$. An identity-covariance Gaussian (right) undergoes a Gaussian diffusion process with learned mean and covariance functions, and is gradually transformed back into the data distribution (left).




Figure 2. Binary sequence learning via binomial diffusion. A binomial diffusion model was trained on binary ‘heartbeat’ data, where a pulse occurs every 5th bin. Generated samples (left) are identical to the training data. The sampling procedure consists of initialization at independent binomial noise (right), which is then transformed into the data distribution by a binomial diffusion process, with trained bit flip probabilities. Each row contains an independent sample. For ease of visualization, all samples have been shifted so that a pulse occurs in the first column. In the raw sequence data, the first pulse is uniformly distributed over the first five bins.

For the experiments shown below, $q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})$ corresponds to either Gaussian diffusion into a Gaussian distribution with identity-covariance, or binomial diffusion into an independent binomial distribution. Table C.1 gives the diffusion kernels for both Gaussian and binomial distributions.

2.2. Reverse Trajectory

The generative distribution will be trained to describe the same trajectory, but in reverse,

$$p(\mathbf{x}^{(T)}) = \pi(\mathbf{x}^{(T)}) \quad (4)$$

$$p(\mathbf{x}^{(0\dots T)}) = p(\mathbf{x}^{(T)}) \prod_{t=1}^T p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}). \quad (5)$$

For both Gaussian and binomial diffusion, for continuous diffusion (limit of small step size β) the reversal of the diffusion process has the identical functional form as the forward process [6]. Since $q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})$ is a Gaussian (binomial) distribution, and if β_t is small, then $q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$



Figure 3. The proposed framework trained on the CIFAR-10 [20] dataset. (a) Example training data. (b) Random samples generated by the diffusion model.

will also be a Gaussian (binomial) distribution. The longer the trajectory the smaller the diffusion rate β can be made.

During learning only the mean and covariance for a Gaussian diffusion kernel, or the bit flip probability for a binomial kernel, need be estimated. As shown in Table C.1, $f_\mu(\mathbf{x}^{(t)}, t)$ and $f_\Sigma(\mathbf{x}^{(t)}, t)$ are functions defining the mean and covariance of the reverse Markov transitions for a Gaussian, and $f_b(\mathbf{x}^{(t)}, t)$ is a function providing the bit flip probability for a binomial distribution. For all results in this paper, multi-layer perceptrons are used to define these functions. A wide range of regression or function fitting techniques would be applicable however, including nonparametric methods.

2.3. Model Probability

The probability the generative model assigns to the data is

$$p(\mathbf{x}^{(0)}) = \int d\mathbf{x}^{(1\cdots T)} p(\mathbf{x}^{(0\cdots T)}). \quad (6)$$

Naively this integral is intractable – but taking a cue from annealed importance sampling and the Jarzynski equality, we instead evaluate the relative probability of the forward and reverse trajectories, averaged over forward trajectories,

$$p(\mathbf{x}^{(0)}) = \int d\mathbf{x}^{(1\cdots T)} p(\mathbf{x}^{(0\cdots T)}) \frac{q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)})}{q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)})} \quad (7)$$

$$= \int d\mathbf{x}^{(1\cdots T)} q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)}) \frac{p(\mathbf{x}^{(0\cdots T)})}{q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)})} \quad (8)$$

$$= \int d\mathbf{x}^{(1\cdots T)} q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)}) \cdot p(\mathbf{x}^{(T)}) \prod_{t=1}^T \frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})}. \quad (9)$$

This can be evaluated rapidly by averaging over samples from the forward trajectory $q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)})$. For infinitesimal β the forward and reverse distribution over trajectories can be made identical (see Section 2.2). If they are identical

then only a *single* sample from $q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)})$ is required to exactly evaluate the above integral, as can be seen by substitution. This corresponds to the case of a quasi-static process in statistical physics [35, 16].

2.4. Training

Training amounts to maximizing the model log likelihood,

$$L = \int d\mathbf{x}^{(0)} q(\mathbf{x}^{(0)}) \log p(\mathbf{x}^{(0)}) \quad (10)$$

$$= \int d\mathbf{x}^{(0)} q(\mathbf{x}^{(0)}).$$

$$\log \left[\frac{\int d\mathbf{x}^{(1\cdots T)} q(\mathbf{x}^{(1\cdots T)}|\mathbf{x}^{(0)}) \cdot p(\mathbf{x}^{(T)}) \prod_{t=1}^T \frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})}}{p(\mathbf{x}^{(0)})} \right], \quad (11)$$

which has a lower bound provided by Jensen’s inequality,

$$L \geq \int d\mathbf{x}^{(0\cdots T)} q(\mathbf{x}^{(0\cdots T)}).$$

$$\log \left[p(\mathbf{x}^{(T)}) \prod_{t=1}^T \frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})} \right]. \quad (12)$$

As described in Appendix B, for our diffusion trajectories this reduces to,

$$L \geq K \quad (13)$$

$$K = - \sum_{t=2}^T \int d\mathbf{x}^{(0)} d\mathbf{x}^{(t)} q(\mathbf{x}^{(0)}, \mathbf{x}^{(t)}) \cdot$$

$$D_{KL} \left(q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}, \mathbf{x}^{(0)}) \parallel p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \right) + H_q(\mathbf{X}^{(T)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(1)}|\mathbf{X}^{(0)}) - H_p(\mathbf{X}^{(T)}). \quad (14)$$

where the entropies and KL divergences can be analytically computed.

As in Section 2.3 if the forward and reverse trajectories are identical, corresponding to a quasi-static process, then the inequality in Equation 13 becomes an equality.

Training consists of finding the reverse Markov transitions which maximize this lower bound on the log likelihood,

$$\hat{p}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) = \operatorname{argmax}_{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})} K. \quad (15)$$

The specific targets of estimation for Gaussian and binomial diffusion are given in Table C.1.

The task of estimating a probability distribution has been reduced to the task of performing regression on the functions which set the mean and covariance of a sequence of Gaussians (or set the state flip probability for a sequence of Bernoulli trials).

2.4.1. SETTING THE DIFFUSION RATE β_t

The choice of β_t in the forward trajectory is important for the performance of the trained model. In AIS, the right schedule of intermediate distributions can greatly improve the accuracy of the log partition function estimate [11]. In thermodynamics the schedule taken when moving between equilibrium distributions determines how much free energy is lost [35, 16].

In the case of Gaussian diffusion, we learn the forward diffusion schedule $\beta_{1\dots T}$ by gradient descent on K . The dependence of samples from $q(\mathbf{x}^{(1\dots T)}|\mathbf{x}^{(0)})$ on $\beta_{1\dots T}$ is made explicit by using ‘frozen noise’ – as in [19] the noise is treated as an additional auxiliary variable, and held constant while computing partial derivatives of K with respect to the parameters.

For binomial diffusion, the discrete state space makes gradient descent with frozen noise impossible². We instead choose the forward diffusion schedule $\beta_{1\dots T}$ to erase a constant fraction $\frac{1}{T}$ of the original signal per diffusion step, yielding a diffusion rate of $\beta_t = (T - t + 1)^{-1}$.

2.5. Multiplying Distributions, and Computing Posteriors

It is common to multiply a model distribution $p(\mathbf{x}^{(0)})$ with a second distribution $r(\mathbf{x}^{(0)})$, producing a new distribution $\tilde{p}(\mathbf{x}^{(0)}) \propto p(\mathbf{x}^{(0)}) r(\mathbf{x}^{(0)})$. This is required by tasks such as computing a posterior in order to do signal denoising or inpainting.

In order to compute $\tilde{p}(\mathbf{x}^{(0)})$, we will multiply each of the intermediate distributions by a corresponding function $r(\mathbf{x}^{(t)})$. A tilde above a distribution or Markov transition indicates that it belongs to a trajectory that has been modified in this way. $\tilde{q}(\mathbf{x}^{(0\dots T)})$ is the modified forward trajectory, which starts at the distribution $\tilde{q}(\mathbf{x}^{(0)}) =$

$\frac{1}{Z_0} q(\mathbf{x}^{(0)}) r(\mathbf{x}^{(0)})$ and proceeds through the sequence of intermediate distributions

$$\tilde{q}(\mathbf{x}^{(t)}) = \frac{1}{\tilde{Z}_t} q(\mathbf{x}^{(t)}) r(\mathbf{x}^{(t)}), \quad (16)$$

where \tilde{Z}_t is the normalizing constant for the t th intermediate distribution. By writing the relationship between the forward and reverse conditional distributions, we can see how this changes the Markov diffusion chain. By Bayes’ rule the forward chain presented in section 2.1 satisfies

$$q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) q(\mathbf{x}^{(t-1)}) = q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) q(\mathbf{x}^{(t)}). \quad (17)$$

The new chain must instead satisfy

$$\tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \tilde{q}(\mathbf{x}^{(t-1)}) = \tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \tilde{q}(\mathbf{x}^{(t)}). \quad (18)$$

As derived in Appendix C, one way to choose a new Markov chain which satisfies Equation 18 is to set

$$\tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \propto q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)}), \quad (19)$$

$$\tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \propto q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}). \quad (20)$$

So that $\tilde{p}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$ corresponds to $\tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$, $p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$ is modified in the corresponding fashion,

$$\tilde{p}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \propto p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}). \quad (21)$$

If $r(\mathbf{x}^{(t-1)})$ is sufficiently smooth, both in \mathbf{x} and t , then it can be treated as a small perturbation to the reverse diffusion kernel $p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$. In this case $\tilde{p}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$ will have an identical functional form to $p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$, but with perturbed mean and covariance for the Gaussian kernel, or with perturbed flip rate for the binomial kernel. The perturbed diffusion kernels are given in Table C.1.

Typically, $r(\mathbf{x}^{(t)})$ should be chosen to change slowly over the course of the trajectory. For the examples in this paper we chose it to be constant,

$$r(\mathbf{x}^{(t)}) = r(\mathbf{x}^{(0)}). \quad (22)$$

Another convenient choice is $r(\mathbf{x}^{(t)}) = r(\mathbf{x}^{(0)})^{\frac{T-t}{T}}$. Under this second choice sampling remains simple even if $r(\mathbf{x}^{(0)})$ is not a tractable distribution, since $r(\mathbf{x}^{(t)})$ makes no contribution at time T , and it thus only enters as a perturbation to the diffusion kernel.

²It would likely be possible to use a Rao-Blackwellization scheme to compute a learning gradient despite the discrete state space.




Figure 4. The proposed framework trained on dead leaf images [24]. (a) Example training image. (b) A sample from the previous state of the art natural image model [38] trained on identical data, reproduced here with permission. (c) A sample generated by the diffusion model. Note that it demonstrates fairly consistent occlusion relationships, displays a multiscale distribution over object sizes, and produces circle-like objects, especially at smaller scales. As shown in Table 2, the diffusion model has the highest log likelihood on the test set.

Dataset	K	$K - L_{null}$
Swiss Roll	2.35 bits	6.45 bits
Binary Heartbeat	-2.414 bits/seq.	12.024 bits/seq.
Bark	-0.55 bits/pixel	1.5 bits/pixel
Dead Leaves	1.489 bits/pixel	3.536 bits/pixel
CIFAR-10	11.895 bits/pixel	18.037 bits/pixel
MNIST		See table 2

Table 1. The lower bound K on the log likelihood, computed on a holdout set, for each of the trained models. See Equation 12. The right column is the improvement relative to an isotropic Gaussian or independent binomial distribution. L_{null} is the log likelihood of $\pi(\mathbf{x}^{(0)})$.

2.6. Entropy of Reverse Process

Since the forward process is known, it is possible to place upper and lower bounds on the entropy of each step in the reverse trajectory. These bounds can be used to constrain the learned reverse transitions $p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$. The bounds on the conditional entropy of a step in the reverse trajectory are

$$\begin{aligned}
 H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}) + H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(0)}) \\
 \leq H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(t)}) \leq H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}),
 \end{aligned}
 \tag{23}$$

where both the upper and lower bounds depend only on the conditional forward trajectory $q(\mathbf{x}^{(1..T)}|\mathbf{x}^{(0)})$, and can be analytically computed. The derivation is provided in Appendix A.

Model	Log Likelihood
Dead Leaves	
MCGSM [38]	1.244 bits/pixel
Diffusion	1.489 bits/pixel
MNIST	
Stacked CAE [1]	121 ± 1.6 bits
DBN [1]	138 ± 2 bits
Deep GSN [2]	214 ± 1.1 bits
Diffusion	218 ± 1.9 bits
Adversarial net [9]	225 ± 2 bits

Table 2. Log likelihood comparisons to other algorithms. Dead leaves images were evaluated using identical training and test data as in [38]. MNIST log likelihoods were estimated using the Parzen-window code from [9], and show that our performance is comparable to other recent techniques.

3. Results

We train diffusion based probabilistic models on several continuous and binary datasets, specifically swiss roll data, binary “heartbeat” data, MNIST [23], CIFAR-10 [20], dead leaf images [24], and bark texture images [22]. We then demonstrate sampling and inpainting of missing data, and compare model log likelihoods and samples to other techniques. In all cases the objective function and gradient were computed using Theano [3], and model training was with SFO [34]. The lower bound on the log likelihood provided by our model is reported for all datasets in Table 1.

3.1. Toy Problems

3.1.1. SWISS ROLL

A probabilistic model was built of a two dimensional swiss roll distribution. The generative model $p(\mathbf{x}^{(0..T)})$ consisted of 40 time steps of Gaussian diffusion initialized




Figure 5. Inpainting. (a) A bark image from [22]. (b) The same image with the central 100×100 pixel region replaced with isotropic Gaussian noise. This is the initialization $\tilde{p}(\mathbf{x}^{(T)})$ for the reverse trajectory. (c) The central 100×100 region has been inpainted using a probabilistic model trained on images of bark. Note the long-range spatial structure, for instance in the crack entering on the left side of the inpainted region.

at an identity-covariance Gaussian distribution. A (normalized) radial basis function network with a single hidden layer and 16 hidden units was trained to generate the mean and covariance functions $\mathbf{f}_\mu(\mathbf{x}^{(t)}, t)$ and a diagonal $\mathbf{f}_\Sigma(\mathbf{x}^{(t)}, t)$ for the reverse trajectory. The top, readout, layer for each function was learned independently for each time step, but for all other layers weights were shared across all time steps and both functions. The top layer output of $\mathbf{f}_\Sigma(\mathbf{x}^{(t)}, t)$ was passed through a sigmoid to restrict it between 0 and 1. As can be seen in Figure 1, the swiss roll distribution was successfully learned.

3.1.2. BINARY HEARTBEAT DISTRIBUTION

A probabilistic model was trained on simple binary sequences of length 20, where a 1 occurs every 5th time bin, and the remainder of the bins are 0. The generative model consisted of 2000 time steps of binomial diffusion initialized at an independent binomial distribution with the same mean activity as the data ($p(x_i^{(T)} = 1) = 0.2$). A multilayer perceptron with sigmoid nonlinearities, 20 input units and three hidden layers with 50 units each was trained to generate the Bernoulli rates $\mathbf{f}_b(\mathbf{x}^{(t)}, t)$ of the reverse trajectory. The top, readout, layer was learned independently for each time step, but for all other layers weights were shared across all time steps. The top layer output was passed through a sigmoid to restrict it between 0 and 1. As can be seen in Figure 2, the heartbeat distribution was successfully learned. The log likelihood under the true generating process is $\log_2(\frac{1}{5}) = -2.322$ bits per sequence. As can be seen in Figure 2 and Table 1 learning was nearly perfect.

3.2. Images

We trained Gaussian diffusion probabilistic models on several image datasets. We begin by describing the architec-

tural components that are shared by the image models. We then present the results for each training set. The architecture is illustrated in Figure 6.

3.2.1. ARCHITECTURE

Readout In all cases, a convolutional network was used to produce a vector of outputs $\mathbf{y}_i \in \mathcal{R}^{2J}$ for each image pixel i . The entries in \mathbf{y}_i are divided into two equal sized subsets, \mathbf{y}^μ and \mathbf{y}^Σ .

Temporal Dependence The convolution output \mathbf{y}^μ is used as per-pixel weighting coefficients in a sum over time-dependent ‘‘bump’’ functions, generating an output $\mathbf{z}_i^\mu \in \mathcal{R}$ for each pixel i ,

$$\mathbf{z}_i^\mu = \sum_{j=1}^J \mathbf{y}_{ij}^\mu g_j(t). \quad (24)$$

The bump functions consist of

$$g_j(t) = \frac{\exp\left(-\frac{1}{2w^2}(t - \tau_j)^2\right)}{\sum_{k=1}^J \exp\left(-\frac{1}{2w^2}(t - \tau_k)^2\right)}, \quad (25)$$

where $\tau_j \in (0, T)$ is the bump center, and w is the spacing between bump centers. \mathbf{z}^Σ is generated in an identical way, but using \mathbf{y}^Σ .

Mean and Variance Finally, these outputs are combined to produce a diffusion mean and variance prediction for each pixel i ,

$$\Sigma_{ii} = \sigma(z_i^\Sigma + \sigma^{-1}(\beta_t)), \quad (26)$$

$$\mu_i = (x_i - z_i^\mu)(1 - \Sigma_{ii}) + z_i^\mu. \quad (27)$$

where both Σ and μ are parameterized as a perturbation around the forward diffusion kernel $T_\pi(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}; \beta_t)$,

and z_i^μ is the mean of the equilibrium distribution that would result from applying $p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})$ many times. Σ is restricted to be a diagonal matrix.

Multi-Scale Convolution We wish to accomplish goals that are often achieved with pooling networks – specifically, we wish to discover and make use of long-range and multi-scale dependencies in the training data. However, since the network output is a vector of coefficients for every pixel it is important to generate a full resolution rather than down-sampled feature map. We therefore define multi-scale-convolution layers that consist of the following steps:

1. Perform mean pooling to downsample the image to multiple scales. Downsampling is performed in powers of two.
2. Performing convolution at each scale.
3. Upsample all scales to full resolution, and sum the resulting images.
4. Perform a pointwise nonlinear transformation, consisting of a soft relu ($\log[1 + \exp(\cdot)]$).

Dense Layers Dense (acting on the full image vector) and kernel-width-1 convolutional (acting separately on the feature vector for each pixel) layers share the same form. They consist of a linear transformation, followed by a tanh nonlinearity.

3.2.2. DATASETS

MNIST In order to allow a direct comparison against previous work on a simple dataset, we trained on MNIST digits [23]. The relative log likelihoods are given in Table 2. Samples from the MNIST model are given in Figure App.1 in the Appendix. Our training algorithm provides an asymptotically exact lower bound on the log likelihood. However, most previous reported results on MNIST log likelihood rely on Parzen-window based estimates computed from model samples. For this comparison we therefore estimate MNIST log likelihood using the Parzen-window code released with [9].

CIFAR-10 A probabilistic model was fit to the training images for the CIFAR-10 challenge dataset [20]. Samples from the trained model are provided in Figure 3.

Dead Leaf Images Dead leaf images [24] consist of layered occluding circles, drawn from a power law distribution over scales. They have an analytically tractable structure, but capture many of the statistical complexities of natural images, and therefore provide a compelling test case for natural image models.




Figure 6. Network architecture for mean function $\mathbf{f}_\mu(\mathbf{x}^{(t)}, t)$ and covariance function $\mathbf{f}_\Sigma(\mathbf{x}^{(t)}, t)$, for experiments in Section 3.2. The input image $\mathbf{x}^{(t)}$ passes through several layers of multi-scale convolution (Section 3.2.1). It then passes through several convolutional layers with 1×1 kernels. This is equivalent to a dense transformation performed on each pixel. A linear transformation generates coefficients for readout of both mean $\mu^{(t)}$ and covariance $\Sigma^{(t)}$ for each pixel. Finally, a time dependent readout function converts those coefficients into mean and covariance images, as described in Section 3.2.1. For CIFAR-10 a dense (or fully connected) pathway was used in parallel to the multi-scale convolutional pathway. For MNIST, the dense pathway was used to the exclusion of the multi-scale convolutional pathway.

A probabilistic model was fit to dead leaf images. Training and testing data was identical to that used in the previous state of the art natural image model[38], and we are thus able to directly compare log likelihood of the learned models in Table 2, where we achieve state of the art performance. Samples are presented in Figure 4.

Bark Texture Images A probabilistic model was trained on bark texture images (T01-T04) from [22]. For this dataset we demonstrate that it is straightforward to evaluate or generate from a posterior distribution, by inpainting a large region of missing data using a sample from the model posterior in Figure 5.

4. Conclusion

We have introduced a novel algorithm for modeling probability distributions that enables exact sampling and evaluation of probabilities and demonstrated its effectiveness on a

variety of toy and real datasets, including challenging natural image datasets. For each of these tests we used a largely unchanged version of the basic algorithm, showing that our method can accurately model a wide variety of distributions with minimal modification. Most existing density estimation techniques must sacrifice modeling power in order to stay tractable and efficient, and sampling or evaluation are often extremely expensive. The core of our algorithm consists of estimating the reversal of a Markov diffusion chain which maps data to a noise distribution; as the number of steps is made large, the reversal distribution of each diffusion step becomes simple and easy to estimate. The result is an algorithm that can learn a fit to any data distribution, but which remains tractable to train, *exactly* sample from, and evaluate.

Acknowledgements

We thank Lucas Theis, Subhaneil Lahiri, Ben Poole, Diederik P. Kingma, and Taco Cohen for extremely helpful discussion, and Ian Goodfellow for sharing Parzen-window code. We thank Khan Academy and the Office of Naval Research for funding Jascha Sohl-Dickstein. We further thank the Office of Naval Research, the Burroughs-Wellcome foundation, Sloan foundation, and James S. McDonnell foundation for funding Surya Ganguli.

References

- [1] Yoshua Bengio, Grégoire Mesnil, Yann Dauphin, and Salah Rifai. Better Mixing via Deep Representations. *arXiv preprint arXiv:1207.4404*, July 2012.
- [2] Yoshua Bengio and Eric Thibodeau-Laufer. Deep generative stochastic networks trainable by backprop. *arXiv preprint arXiv:1306.1091*, 2013.
- [3] J Bergstra and O Breuleux. Theano: a CPU and GPU math expression compiler. *Proceedings of the Python for Scientific Computing Conference (SciPy)*, 2010.
- [4] Julian Besag. Statistical Analysis of Non-Lattice Data. *The Statistician*, 24(3), 179-195, 1975.
- [5] Peter Dayan, Geoffrey E Hinton, Radford M Neal, and Richard S Zemel. The helmholtz machine. *Neural computation*, 7(5):889-904, 1995.
- [6] William Feller and Others. On the theory of stochastic processes, with particular reference to applications. In *Proceedings of the [First] Berkeley Symposium on Mathematical Statistics and Probability*. The Regents of the University of California, 1949.
- [7] Samuel J Gershman and David M Blei. A tutorial on Bayesian nonparametric models. *Journal of Mathematical Psychology*, 56(1):1-12, 2012.
- [8] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359-378, 2007.
- [9] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. *Advances in Neural Information Processing Systems*, 2014.
- [10] Karol Gregor, Ivo Danihelka, Andriy Mnih, Charles Blundell, and Daan Wierstra. Deep AutoRegressive Networks. *arXiv preprint arXiv:1310.8499*, October 2013.
- [11] Roger B Grosse, Chris J Maddison, and Ruslan Salakhutdinov. Annealing between distributions by averaging moments. In *Advances in Neural Information Processing Systems*, pages 2769-2777, 2013.
- [12] G E Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771-1800, 2002.
- [13] Geoffrey E Hinton. The wake-sleep algorithm for unsupervised neural networks). *Science*, 1995.
- [14] A Hyvärinen. Estimation of non-normalized statistical models using score matching. *Journal of Machine Learning Research*, 6:695-709, 2005.
- [15] C Jarzynski. Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach. *Physical Review E*, January 1997.
- [16] Christopher Jarzynski. Equalities and inequalities: irreversibility and the second law of thermodynamics at the nanoscale. In *Annu. Rev. Condens. Matter Phys.* Springer, 2011.
- [17] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183-233, 1999.
- [18] Koray Kavukcuoglu, Marc Aurelio Ranzato, and Yann LeCun. Fast inference in sparse coding algorithms with applications to object recognition. *arXiv preprint arXiv:1010.3467*, 2010.
- [19] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. *International Conference on Learning Representations*, December 2013.
- [20] A Krizhevsky and G Hinton. Learning multiple layers of features from tiny images. *Computer Science Department University of Toronto Tech. Rep.*, 2009.
- [21] Paul Langevin. Sur la théorie du mouvement brownien. *CR Acad. Sci. Paris*, 146(530-533), 1908.
- [22] Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. A sparse texture representation using local affine regions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(8):1265-1278, 2005.
- [23] Y LeCun and C Cortes. The MNIST database of handwritten digits. 1998.
- [24] AB Lee, D Mumford, and J Huang. Occlusion models for natural images: A statistical study of a scale-invariant dead leaves model. *International Journal of Computer Vision*, 2001.
- [25] Siwei Lyu. Unifying Non-Maximum Likelihood Learning Objectives with Minimum KL Contraction. In J Shawe-Taylor, R S Zemel, P Bartlett, F C N Pereira, and K Q Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 64-72. 2011.
- [26] Kevin P Murphy, Yair Weiss, and Michael I Jordan. Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pages 467-475. Morgan Kaufmann Publishers Inc., 1999.
- [27] R Neal. Annealed importance sampling. *Statistics and Computing*, January 2001.
- [28] Sherjil Ozair and Yoshua Bengio. Deep Directed Generative Autoencoders. October 2014.
- [29] Matthew Parry, A Philip Dawid, Steffen Lauritzen, and Others. Proper local scoring rules. *The Annals of Statistics*, 40(1):561-592, 2012.
- [30] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, January 2014.
- [31] Cristian Sminchisescu, Atul Kanaujia, and Dimitris Metaxas. Learning joint top-down and bottom-up processes for 3D visual inference. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1743-1752. IEEE, 2006.
- [32] Jascha Sohl-Dickstein, Peter Battaglino, and Michael DeWeese. New Method for Parameter Estimation in Probabilistic Models: Minimum Probability Flow. *Physical Review Letters*, 107(22):11-14, November 2011.
- [33] Jascha Sohl-Dickstein, Peter B. Battaglino, and Michael R. DeWeese. Minimum Probability Flow Learning. *International Conference on Machine Learning*, 107(22):11-14, November 2011.

- [34] Jascha Sohl-Dickstein, Ben Poole, and Surya Ganguli. Fast large-scale optimization by unifying stochastic gradient and quasi-Newton methods. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 604–612, 2014.
- [35] Richard Spinney and Ian Ford. Fluctuation Relations : A Pedagogical Overview. *arXiv preprint arXiv:1201.6381*, pages 3–56, 2013.
- [36] Plefka T. Convergence condition of the TAP equation for the infinite-ranged Ising spin glass model. *J. Phys. A: Math. Gen.* 15 1971, 1982.
- [37] T Tanaka. Mean-field theory of Boltzmann machine learning. *Physical Review Letters E*, January 1998.
- [38] Lucas Theis, Reshad Hosseini, and Matthias Bethge. Mixtures of conditional Gaussian scale mixtures applied to multiscale image representations. *PLoS one*, 7(7):e39857, 2012.
- [39] M Welling and G Hinton. A new learning algorithm for mean field Boltzmann machines. *Lecture Notes in Computer Science*, January 2002.
- [40] Li Yao, Sherjil Ozair, Kyunghyun Cho, and Yoshua Bengio. On the Equivalence Between Deep NADE and Generative Stochastic Networks. In *Machine Learning and Knowledge Discovery in Databases*, pages 322–336. Springer, 2014.

Appendix

A. Conditional Entropy Bounds Derivation

The conditional entropy $H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(t)})$ of a step in the reverse trajectory is

$$H_q(\mathbf{X}^{(t-1)}, \mathbf{X}^{(t)}) = H_q(\mathbf{X}^{(t)}, \mathbf{X}^{(t-1)}) \quad (28)$$

$$H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(t)}) + H_q(\mathbf{X}^{(t)}) = H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}) + H_q(\mathbf{X}^{(t-1)}) \quad (29)$$

$$H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(t)}) = H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}) + H_q(\mathbf{X}^{(t-1)}) - H_q(\mathbf{X}^{(t)}) \quad (30)$$

An upper bound on the entropy change can be constructed by observing that $\pi(\mathbf{y})$ is the maximum entropy distribution. This holds without qualification for the binomial distribution, and holds for variance 1 training data for the Gaussian case. For the Gaussian case, training data must therefore be scaled to have unit norm for the following equalities to hold. It need not be whitened. The upper bound is derived as follows,

$$H_q(\mathbf{X}^{(t)}) \geq H_q(\mathbf{X}^{(t-1)}) \quad (31)$$

$$H_q(\mathbf{X}^{(t-1)}) - H_q(\mathbf{X}^{(t)}) \leq 0 \quad (32)$$

$$H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(t)}) \leq H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}). \quad (33)$$

A lower bound on the entropy difference can be established by observing that additional steps in a Markov chain do not increase the information available about the initial state in the chain, and thus do not decrease the conditional entropy of the initial state,

$$H_q(\mathbf{X}^{(0)}|\mathbf{X}^{(t)}) \geq H_q(\mathbf{X}^{(0)}|\mathbf{X}^{(t-1)}) \quad (34)$$

$$H_q(\mathbf{X}^{(t-1)}) - H_q(\mathbf{X}^{(t)}) \geq H_q(\mathbf{X}^{(0)}|\mathbf{X}^{(t-1)}) + H_q(\mathbf{X}^{(t-1)}) - H_q(\mathbf{X}^{(0)}|\mathbf{X}^{(t)}) - H_q(\mathbf{X}^{(t)}) \quad (35)$$

$$H_q(\mathbf{X}^{(t-1)}) - H_q(\mathbf{X}^{(t)}) \geq H_q(\mathbf{X}^{(0)}, \mathbf{X}^{(t-1)}) - H_q(\mathbf{X}^{(0)}, \mathbf{X}^{(t)}) \quad (36)$$

$$H_q(\mathbf{X}^{(t-1)}) - H_q(\mathbf{X}^{(t)}) \geq H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(0)}) \quad (37)$$

$$H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(t)}) \geq H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}) + H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(0)}). \quad (38)$$

Combining these expressions, we bound the conditional entropy for a single step,

$$H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}) \geq H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(t)}) \geq H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(t-1)}) + H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(0)}), \quad (39)$$

where both the upper and lower bounds depend only on the conditional forward trajectory $q(\mathbf{x}^{(1 \dots T)}|\mathbf{x}^{(0)})$, and can be analytically computed.

B. Log Likelihood Lower Bound

The lower bound on the log likelihood is

$$L \geq \int d\mathbf{x}^{(0 \dots T)} q(\mathbf{x}^{(0 \dots T)}) \log \left[p(\mathbf{x}^{(T)}) \prod_{t=1}^T \frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})} \right] \quad (40)$$

$$= \int d\mathbf{x}^{(0 \dots T)} q(\mathbf{x}^{(0 \dots T)}) \sum_{t=1}^T \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})} \right] + \int d\mathbf{x}^{(T)} q(\mathbf{x}^{(T)}) \log p(\mathbf{x}^{(T)}). \quad (41)$$

By design, the cross entropy is constant under our diffusion kernels, and equal to the entropy of $p(\mathbf{x}^{(T)})$. Therefore

$$L \geq \sum_{t=1}^T \int d\mathbf{x}^{(0\dots T)} q(\mathbf{x}^{(0\dots T)}) \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})} \right] - H_p(\mathbf{X}^{(T)}). \quad (42)$$

In order to avoid edge effects, we set $p(\mathbf{x}^{(0)}|\mathbf{x}^{(1)}) = q(\mathbf{x}^{(1)}|\mathbf{x}^{(0)}) \frac{p(\mathbf{x}^{(0)})}{p(\mathbf{x}^{(1)})}$, in which case

$$L \geq \sum_{t=2}^T \int d\mathbf{x}^{(0\dots T)} q(\mathbf{x}^{(0\dots T)}) \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})} \right] - H_p(\mathbf{X}^{(T)}) + (H_p(\mathbf{X}^{(1)}) - H_p(\mathbf{X}^{(0)})). \quad (43)$$

Since $(H_p(\mathbf{X}^{(1)}) - H_p(\mathbf{X}^{(0)})) \leq 0$, then

$$L \geq K \quad (44)$$

$$K = \sum_{t=2}^T \int d\mathbf{x}^{(0\dots T)} q(\mathbf{x}^{(0\dots T)}) \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)})} \right] - H_p(\mathbf{X}^{(T)}). \quad (45)$$

Because the forward trajectory is a Markov process,

$$K = \sum_{t=2}^T \int d\mathbf{x}^{(0\dots T)} q(\mathbf{x}^{(0\dots T)}) \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}, \mathbf{x}^{(0)})} \right] - H_p(\mathbf{X}^{(T)}). \quad (46)$$

Using Bayes' rule we can rewrite this in terms of a posterior and marginals from the forward trajectory,

$$K = \sum_{t=2}^T \int d\mathbf{x}^{(0\dots T)} q(\mathbf{x}^{(0\dots T)}) \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}, \mathbf{x}^{(0)})} \frac{q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(0)})}{q(\mathbf{x}^{(t)}|\mathbf{x}^{(0)})} \right] - H_p(\mathbf{X}^{(T)}). \quad (47)$$

We then recognize that several terms are conditional entropies,

$$K = \sum_{t=2}^T \int d\mathbf{x}^{(0\dots T)} q(\mathbf{x}^{(0\dots T)}) \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}, \mathbf{x}^{(0)})} \right] + \sum_{t=2}^T [H_q(\mathbf{X}^{(t)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(t-1)}|\mathbf{X}^{(0)})] - H_p(\mathbf{X}^{(T)}) \quad (48)$$

$$= \sum_{t=2}^T \int d\mathbf{x}^{(0\dots T)} q(\mathbf{x}^{(0\dots T)}) \log \left[\frac{p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)})}{q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}, \mathbf{x}^{(0)})} \right] + H_q(\mathbf{X}^{(T)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(0)}|\mathbf{X}^{(0)}) - H_p(\mathbf{X}^{(T)}). \quad (49)$$

Finally we transform the log ratio of probability distributions into a KL divergence,

$$K = - \sum_{t=2}^T \int d\mathbf{x}^{(0)} d\mathbf{x}^{(t)} q(\mathbf{x}^{(0)}, \mathbf{x}^{(t)}) D_{KL} \left(q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}, \mathbf{x}^{(0)}) \parallel p(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \right) + H_q(\mathbf{X}^{(T)}|\mathbf{X}^{(0)}) - H_q(\mathbf{X}^{(0)}|\mathbf{X}^{(0)}) - H_p(\mathbf{X}^{(T)}). \quad (50)$$

Note that the entropies can be analytically computed, and the KL divergence can be analytically computed given $\mathbf{x}^{(0)}$ and $\mathbf{x}^{(t)}$.

C. Markov Kernel of Perturbed Distribution

In Equations 19 and 20, the perturbed diffusion kernels are set as follows (unlike in the text body, we include the normalization constant)

$$\tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) = \frac{q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)})}{\int d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)})}, \quad (51)$$

$$\tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) = \frac{q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)})}{\int d\mathbf{x}^{(t-1)} q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)})}, \quad (52)$$

or writing them instead in terms of the original transitions,

$$q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) = \frac{\tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \int d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)})}{r(\mathbf{x}^{(t)})}, \quad (53)$$

$$q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) = \frac{\tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \int d\mathbf{x}^{(t-1)} q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)})}{r(\mathbf{x}^{(t-1)})}. \quad (54)$$

Similarly, we write Equation 16 in terms of the original forward distributions,

$$q(\mathbf{x}^{(t)}) = \frac{\tilde{q}(\mathbf{x}^{(t)}) \tilde{Z}_t}{r(\mathbf{x}^{(t)})}. \quad (55)$$

We substitute into Equation 17,

$$q(\mathbf{x}^{(t-1)}) q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) = q(\mathbf{x}^{(t)}) q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}), \quad (56)$$

$$\frac{\tilde{q}(\mathbf{x}^{(t-1)}) \tilde{Z}_{t-1}}{r(\mathbf{x}^{(t-1)})} \frac{\tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \int d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)})}{r(\mathbf{x}^{(t)})} = \frac{\tilde{q}(\mathbf{x}^{(t)}) \tilde{Z}_t}{r(\mathbf{x}^{(t)})} \frac{\tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \int d\mathbf{x}^{(t-1)} q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)})}{r(\mathbf{x}^{(t-1)})}, \quad (57)$$

$$\tilde{q}(\mathbf{x}^{(t-1)}) \tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \tilde{Z}_{t-1} \int d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)}) = \tilde{q}(\mathbf{x}^{(t)}) \tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \tilde{Z}_t \int d\mathbf{x}^{(t-1)} q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}). \quad (58)$$

We then substitute $\tilde{Z}_t = \int d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}) r(\mathbf{x}^{(t)})$,

$$\begin{aligned} \tilde{q}(\mathbf{x}^{(t-1)}) \tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \int d\mathbf{x}^{(t-1)} q(\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t-1)}) \int d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)}) \\ = \tilde{q}(\mathbf{x}^{(t)}) \tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \int d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}) r(\mathbf{x}^{(t)}) \int d\mathbf{x}^{(t-1)} q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}), \end{aligned} \quad (59)$$

$$\begin{aligned} \tilde{q}(\mathbf{x}^{(t-1)}) \tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \int d\mathbf{x}^{(t-1)} d\mathbf{x}^{(t)} q(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) q(\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}) \\ = \tilde{q}(\mathbf{x}^{(t)}) \tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \int d\mathbf{x}^{(t-1)} d\mathbf{x}^{(t)} q(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) q(\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)}), \end{aligned} \quad (60)$$

$$\begin{aligned} \tilde{q}(\mathbf{x}^{(t-1)}) \tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) \int d\mathbf{x}^{(t-1)} d\mathbf{x}^{(t)} q(\mathbf{x}^{(t-1)}, \mathbf{x}^{(t)}) r(\mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}) \\ = \tilde{q}(\mathbf{x}^{(t)}) \tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}) \int d\mathbf{x}^{(t-1)} d\mathbf{x}^{(t)} q(\mathbf{x}^{(t-1)}, \mathbf{x}^{(t)}) r(\mathbf{x}^{(t-1)}) r(\mathbf{x}^{(t)}). \end{aligned} \quad (61)$$

We can now cancel the identical integrals on each side, achieving our goal of showing that the choice of perturbed Markov transitions in Equations 19 and 20 satisfy Equation 18,

$$\tilde{q}(\mathbf{x}^{(t-1)}) \tilde{q}(\mathbf{x}^{(t)}|\mathbf{x}^{(t-1)}) = \tilde{q}(\mathbf{x}^{(t)}) \tilde{q}(\mathbf{x}^{(t-1)}|\mathbf{x}^{(t)}). \quad (62)$$




Figure App.1. Samples from a diffusion probabilistic model trained on MNIST digits.

		<i>Gaussian</i>	<i>Binomial</i>
Well behaved (analytically tractable) distribution	$\pi(\mathbf{x}^{(T)}) =$	$\mathcal{N}(\mathbf{x}^{(T)}; \mathbf{0}, \mathbf{I})$	$\mathcal{B}(\mathbf{x}^{(T)}; 0.5)$
Forward diffusion kernel	$q(\mathbf{x}^{(t)} \mathbf{x}^{(t-1)}) =$	$\mathcal{N}(\mathbf{x}^{(t)}; \mathbf{x}^{(t-1)}\sqrt{1-\beta_t}, \mathbf{I}\beta_t)$	$\mathcal{B}(\mathbf{x}^{(t)}; \mathbf{x}^{(t-1)}(1-\beta_t) + 0.5\beta_t)$
Reverse diffusion kernel	$p(\mathbf{x}^{(t-1)} \mathbf{x}^{(t)}) =$	$\mathcal{N}(\mathbf{x}^{(t-1)}; \mathbf{f}_\mu(\mathbf{x}^{(t)}, t), \mathbf{f}_\Sigma(\mathbf{x}^{(t)}, t))$	$\mathcal{B}(\mathbf{x}^{(t-1)}; \mathbf{f}_b(\mathbf{x}^{(t)}, t))$
Training targets		$\mathbf{f}_\mu(\mathbf{x}^{(t)}, t), \mathbf{f}_\Sigma(\mathbf{x}^{(t)}, t), \beta_{1..T}$	$\mathbf{f}_b(\mathbf{x}^{(t)}, t)$
Forward distribution	$q(\mathbf{x}^{(0..T)}) =$		$q(\mathbf{x}^{(0)}) \prod_{t=1}^T q(\mathbf{x}^{(t)} \mathbf{x}^{(t-1)})$
Reverse distribution	$p(\mathbf{x}^{(0..T)}) =$		$\pi(\mathbf{x}^{(T)}) \prod_{t=1}^T p(\mathbf{x}^{(t-1)} \mathbf{x}^{(t)})$
Log likelihood	$L =$		$\int d\mathbf{x}^{(0)} q(\mathbf{x}^{(0)}) \log p(\mathbf{x}^{(0)})$
Lower bound on log likelihood	$K =$		$\int d\mathbf{x}^{(0..T)} q(\mathbf{x}^{(0..T)}) \log \left[p(\mathbf{x}^{(T)}) \prod_{t=1}^T \frac{p(\mathbf{x}^{(t-1)} \mathbf{x}^{(t)})}{q(\mathbf{x}^{(t)} \mathbf{x}^{(t-1)})} \right]$
Perturbed forward diffusion kernel	$\tilde{q}(\mathbf{x}^{(t)} \mathbf{x}^{(t-1)}) =$	$\mathcal{N}\left(\mathbf{x}^{(t)}; \mathbf{x}^{(t-1)}\sqrt{1-\beta_t} + \left(\frac{\beta_t}{2}\right)^{\frac{1}{2}} \frac{\partial \log r(\mathbf{x}^{(t)})}{\partial \mathbf{x}^{(t)}}, \mathbf{I}\beta_t\right)$	
Perturbed reverse diffusion kernel	$\tilde{q}(\mathbf{x}^{(t)} \mathbf{x}^{(t-1)}) =$	$\mathcal{N}\left(\mathbf{x}^{(t-1)}; \mathbf{f}_\mu(\mathbf{x}^{(t)}, t) + \left(\frac{\mathbf{f}_\Sigma(\mathbf{x}^{(t)}, t)}{2}\right)^{\frac{1}{2}} \frac{\partial \log r(\mathbf{x}^{(t)})}{\partial \mathbf{x}^{(t)}}, \mathbf{f}_\Sigma(\mathbf{x}^{(t)}, t)\right)$	

Table C.1. The key equations in this paper for the specific cases of Gaussian and binomial diffusion processes. $\mathcal{N}(u; \mu, \Sigma)$ is a Gaussian distribution with mean μ and covariance Σ . $\mathcal{B}(u; r)$ is the distribution for a single Bernoulli trial, with $u = 1$ occurring with probability r , and $u = 0$ occurring with probability $1 - r$.