

DIFFUSION MODELS ARE KELLY GAMBLERS

Anonymous authors

Paper under double-blind review

ABSTRACT

We draw a connection between diffusion models and the Kelly criterion for maximizing returns in betting games. A signal that is correlated with the outcome of such a game can be used to focus the bets on a narrow range of high probability predictions. Diffusion models share the same paradigm in that they gradually concentrate the probability mass to fit the training data. We show that the information stored in an unconditional diffusion model captures, in part, the joint correlation between the components of the data variable X . Conditional diffusion models store additional information to bind the signal X with the conditioning information Y , equal to the mutual information between them. The latter is only a small fraction of the total information in the neural network if the data is low-dimensional. We examine why this does not hinder conditional generation.

1 INTRODUCTION

Diffusion models are highly effective at approximating continuous high-dimensional probability distributions like images, audio, and video (Sohl-Dickstein et al., 2015; Song et al., 2021b; Dhariwal & Nichol, 2021), and more recently, discrete data like language (Lou et al., 2024; Nie et al., 2025). They generate samples by progressively denoising random vectors using information gathered during training—information that was eroded by the forward diffusion process as it transformed the data into noise. A method to estimate this information was introduced in Premkumar (2025).

The Kelly criterion is a rule for allocating capital in a betting game when we believe we have an information edge over the odds. In a game with fair odds, our returns are maximized if we place our bets in proportion to the probability of outcomes Kelly (1956). Therefore, Kelly gambling is fundamentally a problem of probabilistic modeling: our winnings depend on how accurately our model captures the true distribution of outcomes. This is also the goal of generative modeling.

The accuracy of our predictions improves if we have access to side information, since it narrows the range of possible outcomes. In particular, the financial value of such information depends on how well it correlates with the game’s outcomes. This is quantified by the mutual information between the outcomes and the side information. The mutual information between two random variables tells us how much knowing one variable reduces our uncertainty about the other, regardless of how complex their relationship is (Shannon, 1948). If the side information is a good indicator of the outcome, the mutual information between them is high, and so are our chances of winning.

The analogous problem in generative modeling is to produce samples of a random variable X , given some training data. A diffusion model that generates samples of X can be used to make bets in a hypothetical game where we wager on the value of X . The bets are Kelly-optimal iff the model captures $p(x)$ perfectly. To do this, a diffusion model stores some information, as measured by the *neural entropy* S_{NN}^x . A conditional model of $X|Y$ can make more targeted bets than an unconditional one since, on average, $p(x|y)$ is narrower than $p(x)$. To remain Kelly-optimal, that is, to reconstruct $p(x|y)$ faithfully, the model must store an additional $I(X; Y)$ nats of information.

Real-world diffusion models are often applied to problems where the mutual information between the data variable \mathbf{X} and the conditioning variable \mathbf{Y} is low compared to the information stored in a model trained on \mathbf{X} unconditionally (see Fig. 1). The latter is quantified by S_{NN}^x , which is different from $S(\mathbf{X})$, the entropy of \mathbf{X} . Neural entropy is very large if \mathbf{X} inhabits a lower-dimensional manifold relative to its naive co-ordinate representation, since it measures the effort required to collapse a Gaussian ball from ambient space to that manifold. This requires a far greater reduction in uncertainty than it would take if \mathbf{X} fully occupied the ambient dimensions. A key point in this

paper is that *the information required to locate the \mathbf{X} manifold precisely is largely irrelevant in correlating \mathbf{X} with \mathbf{Y}* , so the models still produce reasonable estimates of $I(\mathbf{X}; \mathbf{Y})$.

At an operational level, the score function must become large at the later stages of the reverse process to squeeze the probability mass into a thin sliver of the total volume. This is what makes $S_{\text{NN}}^{\mathbf{X}}$ large (cf. Eq. (12)). However, diffusion models can be used to reliably compute $I(\mathbf{X}; \mathbf{Y})$ even in this situation, using a formula that involves the difference of the conditional and unconditional scores, since the divergent behavior cancels in that vector (Franzese et al., 2024; Kong et al., 2024). This is the same as the ‘classification’ vector used to amplify the conditional signal in classifier-free guidance (CFG) (Ho & Salimans, 2022). This is partly why CFG is effective. We show that CFG increases $I(\mathbf{X}; \mathbf{Y})$, albeit at the expense of distorting the \mathbf{X} distribution. Furthermore, the effect of CFG saturates beyond a certain value of the guidance strength parameter.

In image diffusion models, a sizable portion of $S_{\text{NN}}^{\mathbf{X}}$ is used up in resolving the finer texture of the images \mathbf{X} , since reproducing these details is tantamount to locating the manifold precisely (see Sec. 5). We show this by identifying within $S_{\text{NN}}^{\mathbf{X}}$ a term called the *total correlation*, $\text{TC}(\mathbf{X})$, which is a higher-dimensional generalization of mutual information (Watanabe, 1960). It measures the joint correlation between the pixels in \mathbf{X} . To get the textures right, the model must establish tight correlations between neighboring pixels, which can be seen as a multivariate Kelly bet. The small-scale details are largely shared between different image classes and contribute little to distinguishing images with different macroscopic features. This is why $I(\mathbf{X}; \mathbf{Y})$ is low and $S_{\text{NN}}^{\mathbf{X}}$ is large. These properties become manifest if we allow a diffusion model to generate its own conditioning signal \mathbf{Z} by pairing it with a $\mathbf{X} \rightarrow \mathbf{Z}$ encoder. By varying the time intervals that modulate \mathbf{Z} , we show that diffusion models pick up little to no class-specific information from the textures, although they eat up a large fraction of the information budget.

Contributions In this paper we (1) introduce a novel perspective on diffusion models, connecting generative modeling to optimal betting, (2) quantify the amount of information required to correlate the components within \mathbf{X} , and separately \mathbf{X} with \mathbf{Y} , (3) show that the former overwhelms the latter in image diffusion models, $S_{\text{NN}}^{\mathbf{X}} \gg I(\mathbf{X}; \mathbf{Y})$, (4) connect the manifold structure of images \mathbf{X} to their perceptual and semantic content, (5) study how these components correlate with the class information, and (6) demonstrate why $I(\mathbf{X}; \mathbf{Y})$ can be captured despite high $S_{\text{NN}}^{\mathbf{X}}$. **(Right)** Plot showing how $I(\mathbf{X}; \mathbf{Y})$ becomes a vanishingly smaller fraction of the neural entropy $S_{\text{NN}}^{\mathbf{X}}$ as R_{eff}^2 , the effective size of $p(\mathbf{x})$, is lowered. See Sec. 6 for more details.

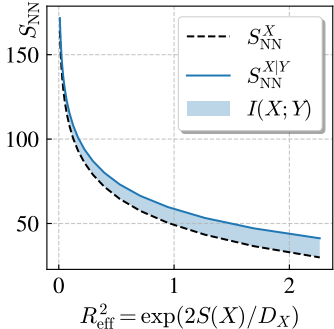


Figure 1

2 KELLY CRITERION

We begin with a brief discussion of Kelly’s argument using a toy example. Simplifying assumptions are made to emphasize ideas central to the rest of the paper. For a more rigorous treatment, see Kelly (1956), Chapter 4 of Cover & Thomas (2006) or Thorp (2011).

Consider a simple betting game involving a cup that conceals a six-sided die. A dealer rattles the die with a firm shake of the cup and invites us to place bets on the outcome. Believing that the die is fair, the dealer offers 6-for-1 odds—if we bet \$1 on a number and win, we receive \$6, and nothing otherwise. Let the random variable X represent the outcome of a single roll. Over n i.i.d. throws, the outcome (x_1, x_2, \dots, x_n) is almost certain to belong to the typical set of $2^{nH(X)}$ members, each having nearly the same probability when n is large. This is the asymptotic equipartition principle (see Ch. 4 of MacKay (2002)). Here $H(X) := \mathbb{E}[-\log_2 p(x)]$ is the Shannon entropy of X . To optimize our gains, we can distribute our seed money V_0 equally over the typical set and let the game run.¹ The winning sequence will return 6^n times the investment placed on it, so we make a total of $V_n = 2^{n(\log_2 6 - H(X))} V_0$ by the end. We profit because the dealer assumed the die was

¹This is equivalent to placing bets sequentially, in proportion to $p(x)$, after each outcome is revealed; see Example 6.3.1 in Cover & Thomas (2006).

108 fair, hence the 6-for-1 odds, while we placed our bets according to the true distribution of outcomes,
 109 $p(x)$. By doing so, we concentrated our bets on a smaller subset of possible outcomes, $2^{nH(X)}$,
 110 which gave us an advantage over the dealer who thought all 6^n outcomes were equally likely. This
 111 is the Kelly criterion: optimal returns are achieved by allocating to each outcome x a portion of the
 112 capital proportional to $p(x)$. Therefore, *our win rate is directly tied to how well we can model $p(x)$* .

113 Now suppose that, unbeknownst to the dealer, there is a communication channel that transmits the
 114 outcome to us before the cup is lifted. If the channel is noiseless, we are guaranteed to win every
 115 round, and we can grow our wealth by a factor of 6^n over n throws by wagering our entire stake on
 116 the correct outcome each time. In other words, a perfect channel eliminates all uncertainty about the
 117 outcome, which is equivalent to setting $H(X) \rightarrow 0$ in the previous analysis.

118 On the other hand, if the channel is noisy, we can no longer be certain that the information we receive
 119 accurately reflects the outcome under the cup; the noise in the channel reintroduces randomness into
 120 the game. However, the side information may be correlated with the outcome if the noise is not too
 121 severe, so we can do better than when we did not know the outcome at all. If Y is the signal received
 122 from the channel, we can place bets according to $p(x|y)$ rather than $p(x)$. That is, we repeat the
 123 arguments from above with $H(X)$ replaced by the conditional entropy $H(X|Y)$, which quantifies
 124 the residual uncertainty about X given access to Y —knowing Y allows us to concentrate our bets on
 125 just $2^{nH(X|Y)}$ sequences. Then, optimal betting yields a final wealth of $V'_n = 2^{n(\log_2 6 - H(X|Y))} V_0$.
 126 Thus, the increase in doubling rate *due to the communication channel* is the mutual information
 127 between X and Y (see App. A.1),

$$128 R := \lim_{n \rightarrow 0} \frac{1}{n} \log_2 \frac{V'_n}{V_n} = H(X) - H(X|Y) \equiv I(X; Y). \quad (1)$$

131 3 THE DIFFUSION GAMBLER

132 Given a set of data vectors $\{\mathbf{x}^{(i)}\}_{i=1}^N$ in $\mathbb{R}^{D_{\mathbf{x}}}$, a probabilistic model approximates the underlying
 133 distribution p_d from which these vectors could have been sampled. One way to do this is to transform
 134 a generic initial distribution p_0 into one that is more likely to have produced the given samples. If
 135 p_0 is nearly the equilibrium state of a diffusive process (see App. G for notation)
 136 p_0 is nearly the equilibrium state of a diffusive process (see App. G for notation)
 137

$$138 d\tilde{\mathbf{X}}_s = b_+(\tilde{\mathbf{X}}_s, s)ds + \sigma(s)d\hat{\mathbf{B}}_s, \quad (2)$$

139 then the transformation we seek is simply a reversal (playback) of the forward evolution that converts
 140 $p_d \rightarrow p_0$ according to Eq. (2). The reverse process is effected by

$$141 d\mathbf{X}_t = -(b_+(\mathbf{X}_t, T-t) - \sigma(T-t)^2 \nabla \log p(\mathbf{X}_t, t))dt + \sigma(T-t)d\mathbf{B}_t, \quad (3)$$

142 where $t := T - s$ is a time variable that runs in the opposite direction to s , and p is the density that
 143 interpolates p_0 and p_d (see Fig. 17). Diffusion is a dissipative process that erases information over
 144 time, which means reversal must reinstate the same amount of information to drive p_0 back to p_d .
 145 If p_d is subject to Eq. (2) for a time T , and b_+, σ have the same time-dependence, the information
 146 that must be injected to return to p_d is quantified by the total entropy produced (Vaikuntanathan &
 147 Jarzynski, 2009; Seifert, 2012),
 148

$$149 S_{\text{tot}} := \int_0^T dt \frac{\sigma^2}{2} \mathbb{E}_p \left[\left\| \nabla \log p_{\text{eq}}^{(t)} - \nabla \log p \right\|^2 \right] = D_{\text{KL}}(p_d \| p_{\text{eq}}^{(T)}) - D_{\text{KL}}(p_0 \| p_{\text{eq}}^{(0)}). \quad (4)$$

150 The expectation is taken over trajectories generated by Eq. (2), starting at $\tilde{\mathbf{X}}_0 \sim p_d$, and $p_{\text{eq}}^{(t)}$ is the
 151 quasi-invariant state, which can be understood as the ‘least informative state’ at time t . It is the
 152 distribution that would result if we froze b_+ and σ at their values at t and waited for the system to
 153 equilibrate. That is, $p_{\text{eq}}^{(t)}(x) \propto \exp \left[\int^x 2b_+ / \sigma^2 \right]$. In a diffusion model, the drift term in Eq. (3) is
 154 approximated by a neural network. It is useful to parameterize the reverse SDE as
 155

$$156 d\mathbf{X}_t = (b_+(\mathbf{X}_t, T-t) + \sigma(T-t)^2 \mathbf{e}_\theta(\mathbf{X}_t, T-t))dt + \sigma(T-t)d\mathbf{B}_t, \quad (5)$$

157 where the neural network \mathbf{e}_θ is trained to minimize (cf. App. B.3)

$$158 \mathcal{L}_{\text{EM}} = \int_0^T dt \frac{\sigma^2}{2} \mathbb{E}_p \left[\left\| \nabla \log p_{\text{eq}}^{(t)} - \nabla \log p + \mathbf{e}_\theta \right\|^2 \right]. \quad (6)$$

If $e_\theta = 0$, Eq. (5) reduces to the forward dynamics, Eq. (2). Let $\mathcal{P}[p_d]$ be the probability that N random vectors from p_0 would be distributed as p_d at $t = T$ under Eq. (2). A *perfectly* trained diffusion model, with the idealized network $e_\theta^* = -2b_+/\sigma^2 + \nabla \log p$, modifies the dynamics to Eq. (3), which is guaranteed to take $p_0 \rightarrow p_d$. Such a network stores/applies precisely S_{tot} worth of information to affect this transformation, since $S_{\text{tot}} = -\frac{1}{N} \log \mathcal{P}[p_d]$. This is why they are called *entropy-matching* models (Premkumar, 2025).

Information negates uncertainty. The idealized entropy-matching model applies S_{tot} worth of information to reconstitute p_d from p_0 in time T . If T is large enough that $p_0 \approx p_{\text{eq}}^{(0)}$, the total entropy can be written as

$$S_{\text{tot}}^{\mathbf{X}} = D_{\text{KL}} \left(p_d(\mathbf{x}) \parallel p_{\text{eq}}^{(T)}(\mathbf{x}) \right) = -S(\mathbf{X}) - \int d\mathbf{x} p_d(\mathbf{x}) \log p_{\text{eq}}^{(T)}(\mathbf{x}). \quad (7)$$

We have introduced a superscript \mathbf{X} in $S_{\text{tot}}^{\mathbf{X}}$ to specify explicitly the random variable whose distribution is being modeled. Crucially, $S_{\text{tot}}^{\mathbf{X}}$ is *not* the same as $S(\mathbf{X})$. In fact, $S_{\text{tot}}^{\mathbf{X}}$ is larger if $S(\mathbf{X})$ is smaller, since the diffusion model must apply more information to locate a narrower $p_d(\mathbf{x})$. This is especially a problem in continuous diffusion models, where $S(\mathbf{X})$ can be arbitrarily small; the differential entropy of a Dirac delta function is $-\infty$. Such divergences arise when \mathbf{X} lives in a lower-dimensional manifold, as we explore further in Sec. 5.

Eq. (7) also reveals an interesting fact about correlations within the components of \mathbf{X} . In the unconditional case, with T large enough that $p_0 \approx p_{\text{eq}}^{(0)}$, the total entropy can be factorized as

$$S_{\text{tot}} = D_{\text{KL}} \left(p_d \parallel p_{\text{eq}}^{(T)} \right) = \sum_{k=1}^{D_{\mathbf{X}}} D_{\text{KL}} \left(p_d(x_k) \parallel p_{\text{eq}}^{(T)}(x_k) \right) + \underbrace{D_{\text{KL}} \left(p_d(x_1, \dots, x_{D_{\mathbf{X}}}) \parallel \prod_{k=1}^{D_{\mathbf{X}}} p_d(x_k) \right)}_{\text{TC}(\mathbf{X})}. \quad (8)$$

The last term, called the *total correlation*, is a generalization of mutual information to multiple random variables (Watanabe, 1960). In Eq. (8) x_k is the k -th component of a data vector \mathbf{x} , and $p_d(x_k)$ and $p_{\text{eq}}^{(T)}(x_k)$ are the marginal densities obtained by integrating $p_d(\mathbf{x})$ and $p_{\text{eq}}^{(T)}(\mathbf{x})$ over all components except x_k . Eq. (8) tells us that during the reversal/generative stage, the model must (1) shift the marginals for each x_k from $p_{\text{eq}}^{(T)}(x_k) \rightarrow p_d(x_k)$, and (2) establish correlations between different x_k . Thus, denoising a vector from p_0 is, in part, the process of *restoring the component-wise correlations* that were lost in the forward stage.

Next, we consider a scenario where the diffusion model is used for conditional generation. Let \mathbf{Y} be the conditioning information. For example, \mathbf{Y} represents the class labels in class-conditioned image generation, with \mathbf{X} being the associated images. Given $\mathbf{Y} = \mathbf{y}$, a new sample can be generated by applying Eq. (5) with

$$e_\theta^*(\mathbf{x}_t, T-t; \mathbf{y}) = -\frac{2b_+(\mathbf{x}_t, T-t)}{\sigma^2(T-t)} + \nabla \log p(\mathbf{x}_t, t | \mathbf{y}). \quad (9)$$

Let $S_{\text{tot}}^{\mathbf{X}|\mathbf{y}}$ denote the information stored by such a network for each \mathbf{y} . On average, this model injects an amount of information

$$S_{\text{tot}}^{\mathbf{X}|\mathbf{Y}} := \mathbb{E}_{\mathbf{Y}} \left[S_{\text{tot}}^{\mathbf{X}|\mathbf{y}} \right] = \mathbb{E}_{\mathbf{Y}} \left[D_{\text{KL}} \left(p_d(\mathbf{x}|\mathbf{y}) \parallel p_{\text{eq}}^{(T)}(\mathbf{x}) \right) \right] = -S(\mathbf{X}|\mathbf{Y}) - \int d\mathbf{x} p_d(\mathbf{x}) \log p_{\text{eq}}^{(T)}(\mathbf{x}). \quad (10)$$

We expect $S_{\text{tot}}^{\mathbf{X}|\mathbf{Y}} \geq S_{\text{tot}}^{\mathbf{X}}$, since more information is needed to squeeze the quasi-invariant state into the distributions $p_d(\mathbf{x}|\mathbf{y})$, which are on average narrower than the marginal $p_d(\mathbf{x})$ (see App. A.1). Indeed, the conditional model injects an *additional* R nats of information, where

$$R := S_{\text{tot}}^{\mathbf{X}|\mathbf{Y}} - S_{\text{tot}}^{\mathbf{X}} = S(\mathbf{X}) - S(\mathbf{X}|\mathbf{Y}) \equiv I(\mathbf{X}; \mathbf{Y}). \quad (11)$$

Importantly, this $I(\mathbf{X}; \mathbf{Y})$ amount of information is stored atop $S_{\text{tot}}^{\mathbf{X}}$, which is different from $S(\mathbf{X})$. Comparing Eqs. (1) and (11), we can view the generative step as a *Kelly bet on the value of \mathbf{X} using the side information \mathbf{Y}* . That is, in a betting game where we must guess the value of \mathbf{X} , a conditional model of $\mathbf{X}|\mathbf{Y}$ can grow our capital at an exponential rate that is $I(\mathbf{X}; \mathbf{Y})$ higher than an unconditional one. This is because the diffusion model stores/injects $I(\mathbf{X}; \mathbf{Y})$ nats of extra information, which it uses to correlate \mathbf{X} to \mathbf{Y} . In a similar way, the $\text{TC}(\mathbf{X})$ component of $S_{\text{tot}}^{\mathbf{X}}$ is used to correlate the components of \mathbf{X} with one another jointly. This can be seen as a multivariate Kelly bet, with the value of each x_k informing the value of all others, and vice versa.

Entropy-matching In the discussion above, we have parameterized the reverse drift in Eq. (5) as $b_+ + \sigma^2 \mathbf{e}_\theta$, which is different from the score-matching parameterization of $-b_+ + \sigma^2 \mathbf{s}_\theta$. The latter forces the network to retain additional information to counteract the repulsive $-b_+$ term, as explained in Premkumar (2025). A simple thought experiment reveals the problem: suppose we take $p_d = p_0 \approx p_{\text{eq}}^{(0)}$. The forward process has little effect on the distribution since p_d is already close to equilibrium. However, the scores for this transformation are still non-zero over the support of $p_{\text{eq}}^{(0)}$, which means the network in a score-matching model must retain information to convert a distribution *back to itself*. On the other hand, an entropy-matching network would store no information in this scenario, as expected. In this sense, entropy-matching makes transparent the correspondence between the network’s information content and the entropy of the underlying data. We can convert a score-matching model to an entropy matching one with the simple substitution $\mathbf{s}_\theta = \nabla \log p_{\text{eq}}^{(t)} + \mathbf{e}_\theta$.

Neural Entropy We derived Eq. (11) under the assumption of an ideal entropy-matching model \mathbf{e}_θ^* , which absorbs exactly S_{tot} units of information during training. In practice, no model achieves this ideal because of the finite number of training epochs, limited batch size, and finite data. However, Premkumar (2025) demonstrates that the amount of information stored in a real network \mathbf{e}_θ is measured through its *neural entropy*,

$$S_{\text{NN}}^{\mathbf{X}} := \int_0^T ds \frac{\sigma(s)^2}{2} \mathbb{E}_p \left[\|\mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s)\|^2 \right] \approx S_{\text{tot}}^{\mathbf{X}}. \quad (12)$$

Notice that setting $\mathbf{e}_\theta \rightarrow \mathbf{e}_\theta^*$ turns $S_{\text{NN}} \rightarrow S_{\text{tot}}$, which follows from Eqs. (4) and (6). Away from this theoretical limit, the neural entropy can be either smaller or larger than the true S_{tot} . For example, when the dataset is sparse, the diffusion model tends to concentrate probability mass around the available samples, demanding greater effort from the network than it requires to reconstitute the true p_d , which may have a more distributed support. Another possibility is that training is not long enough for the network to absorb all of S_{tot} , so neural entropy trails the true value. Nevertheless, with sufficient training and a large enough dataset, Eq. (12) provides a close approximation to S_{tot} (see Fig. 10). These are the entropy-matching models we discuss henceforth.

4 MUTUAL INFORMATION AND GUIDANCE

It is possible to use Eq. (11) to estimate the mutual information between two high-dimensional random variables. Given a set of pairs $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^N$ we can train an entropy-matching model to reconstruct the distribution of \mathbf{X} given \mathbf{Y} and, separately, the marginal distribution of \mathbf{X} . The difference in neural entropy between the two approximates $I(\mathbf{X}; \mathbf{Y})$. This approach is closely related to the results in Franzese et al. (2024). Specifically, their Eq. (19) says

$$I(\mathbf{X}; \mathbf{Y}) = \mathbb{E}_{\mathbf{Y}} \left[\int_0^T ds \frac{\sigma^2}{2} \mathbb{E}_{\tilde{\mathbf{x}}_s, \mathbf{X}|\mathbf{y}} \left[\|\nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y}) - \nabla \log p(\tilde{\mathbf{x}}_s, s)\|^2 \right] \right] \quad (13)$$

$$\approx \mathbb{E}_{\mathbf{Y}} \left[\int_0^T ds \frac{\sigma^2}{2} \mathbb{E}_{\tilde{\mathbf{x}}_s, \mathbf{X}|\mathbf{y}} \left[\|\mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s; \mathbf{y}) - \mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s)\|^2 \right] \right], \quad (14)$$

up to terms that vanish as $T \rightarrow 0$. We give an alternative derivation of this result in App. D.1. Notice that Eq. (14) also works with score-matching models, with \mathbf{e}_θ replaced by the corresponding \mathbf{s}_θ . The main advantage of entropy-matching is that it links the information stored in the network to the effort required to reconstitute p_d . For now, we consider how Eq. (13) helps us better understand classifier-free guidance (CFG) (Ho & Salimans, 2022).

In image models, \mathbf{X} denotes images and \mathbf{Y} the corresponding class labels. Dhariwal & Nichol (2021) show that the quality of generated samples can be improved, at the cost of decreased diversity, by forcing the model to adhere more strongly to the conditioning variable. If Eq. (3) evolves p_0 to $p(\mathbf{x}_t, t|\mathbf{y})$ at time t , the conditioning on \mathbf{y} can be amplified by modifying the drift vectors to sample from $p(\mathbf{x}_t, t|\mathbf{y})p(\mathbf{y}|\mathbf{x}_t, t)^w$ instead, where $w > 0$ is a parameter we can control. Ho & Salimans (2022) accomplish this by constructing an implicit classifier $p(\mathbf{y}|\mathbf{x}_t, t) \propto p(\mathbf{x}_t, t|\mathbf{y})/p(\mathbf{x}_t, t)$, which has the score

$$\begin{aligned} \mathbf{s}_{\text{cl}}(\mathbf{y}, t) &:= \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t, t) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t, t|\mathbf{y}) - \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t, t) \\ &\approx \nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y}) - \nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s) \approx \mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s; \mathbf{y}) - \mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s). \end{aligned} \quad (15)$$

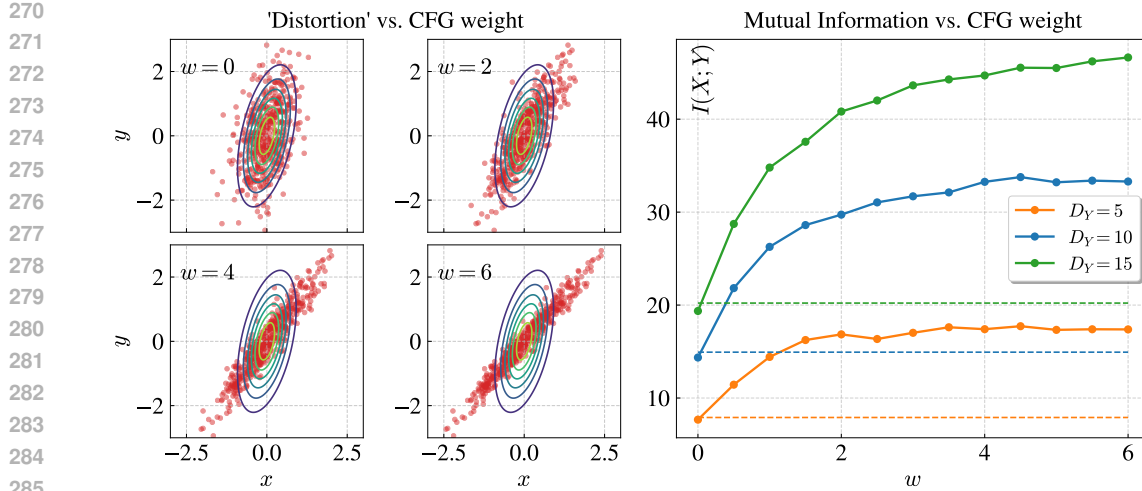


Figure 2: **Left:** Samples generated by a CFG-style modification to the conditional score $\nabla \log p(x_t, t|y)$ of a joint Gaussian, $\mathbf{Y} = \mathbf{A}\mathbf{X} + \varepsilon$ (cf. Eqs. (16) and (53)). CFG strengthens the correlation between \mathbf{X} and \mathbf{Y} , increasing their mutual information. But it also alters the relationship between them. **Right:** Mutual information under CFG for the joint Gaussian. We fix $D_{\mathbf{X}} = 25$ and repeat the experiment with $D_{\mathbf{Y}} = 5, 10, 15$. Notice how $I(\mathbf{X}; \mathbf{Y})_{\text{CFG}}$ increases as the guidance strength is ramped up (dashed lines are the original $I(\mathbf{X}; \mathbf{Y})$ values, cf. Eq. (55)). It saturates faster for smaller $D_{\mathbf{Y}}$, when \mathbf{Y} has fewer degrees of freedom to encode the diversity in \mathbf{X} . See Figs. 7 and 8 for more details.

This is also the vector whose ℓ_2 -norm appears in Eq. (13). A close examination of the latter helps us build some intuition for s_{cl} . We begin by noting that the inner expectation in Eq. (13) is computed over trajectories that start from samples of $\mathbf{X}|\mathbf{Y} = \mathbf{y}$ (or $\mathbf{X}|\mathbf{y}$ for short). If we were to average $\|s_{\text{cl}}(\mathbf{y}, s)\|^2$ over forward paths that emanate from $\mathbf{X}|\mathbf{y}'$ where $\mathbf{y}' \neq \mathbf{y}$, the expectation value would be higher, since the conditional score $\nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y})$ along those paths are larger. That is, the conditional scores are steeper in regions explored by diffused versions of $\mathbf{X}|\mathbf{y}'$, where the underlying conditional density $p(\tilde{\mathbf{x}}_s, s|\mathbf{y})$ falls off rapidly. Therefore, the integral over $\mathbb{E}_{\tilde{\mathbf{x}}_s, \mathbf{X}|\mathbf{y}'}$ $\|s_{\text{cl}}(\mathbf{y}, s)\|^2$ is *smallest* when $\mathbf{y}' = \mathbf{y}$.² During the generative process, a sample can be made more \mathbf{y} -like by augmenting the reverse drift term in Eq. (3) with $w \times s_{\text{cl}}(\mathbf{y}, t)$:

$$\nabla \log p(\mathbf{x}_t, t|\mathbf{y}) \rightarrow (1 + w)\nabla \log p(\mathbf{x}_t, t|\mathbf{y}) - w\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t, t). \quad (16)$$

This is CFG. The new drift term pulls a reverse-evolving \mathbf{x}_t more strongly toward regions in \mathbf{x} -space that are consistent with the condition \mathbf{y} . At the next time step, $\mathbf{x}_{t+\Delta t}$ is more \mathbf{y} -like, so $s_{\text{cl}}(\mathbf{y}, t)$ weakens. Therefore, the effect of CFG diminishes over time t . In the limit, the fully denoised sample will be more tightly determined by \mathbf{y} , so $I(\mathbf{X}; \mathbf{Y})_{\text{CFG}}$ will be higher. There is a limit, however, since the increase in $I(\mathbf{X}; \mathbf{Y})_{\text{CFG}}$ saturates at higher values of w if the dimensionality of \mathbf{Y} is smaller than that of \mathbf{X} (see Fig. 2). In that case, \mathbf{Y} lacks the sufficient code length to encode the information content of \mathbf{X} faithfully; this is an example of an *information bottleneck* (see App. A.2).

One might wonder whether it is possible to substitute Eq. (16) into Eq. (13) to conclude that mutual information is boosted to $(1 + w)^2 I(\mathbf{X}; \mathbf{Y})$. But that would be incorrect; such a maneuver is disallowed by the fact that the modified score does not correspond to any known forward diffusion process. For the same reason, the distribution we reconstruct under CFG is not the true p_d (see Fig. 2). We discuss this point further in App. D.2. It remains true, however, that CFG strengthens the binding between \mathbf{X} and \mathbf{Y} . This is especially useful when the base training cannot reliably ensure that conditioning is respected. The underlying reason for this is a combination of the manifold structure of the training data itself, and our choice to model it with a continuous random variable.

²In fact, it is possible to build a classifier based on this very insight (Clark & Jaini, 2023; Li et al., 2023).

5 MANIFOLDS AND THE INFORMATION BUDGET

Mutual information is largest if knowledge of the value of one variable completely determines the value of the other (see App. A.1). On the other hand, if a given value of \mathbf{Y} corresponds to a wide range of \mathbf{X} , a greater amount of uncertainty remains about the value of the latter, so $I(\mathbf{X}; \mathbf{Y})$ is low. This is the case with labeled image datasets, where a label $\mathbf{Y} = \mathbf{y}$ can correspond to a rich distribution of images $\mathbf{X} | \mathbf{Y} = \mathbf{y}$. For instance, the label ‘dog’ corresponds to a wide variety of dogs. There is, however, a subtle point here about image datasets: the entropy of images does not arise only from high-level semantic variation (different breeds, poses, or scenes), but is overwhelmingly dominated by the low-level perceptual details present in each image. Diffusion models capture these details with remarkable fidelity, and a large share of their information capacity is devoted to encoding fine perceptual structure rather than high-level semantics. In fact, much of the low-level detail is not class-specific, but is shared across multiple categories of images (see Fig. 15). Therefore, specifying the label ‘dog’ does very little to narrow down the possibilities of which sample to draw. This is why diffusion models sometimes stray from the conditioning signal during generation: the mutual information between images and labels is intrinsically low compared to its neural entropy. We provide empirical proof of these statements in Sec. 6. For now, we will rationalize them from an information theory perspective.

We have made two assertions thus far: (1) neural entropy is high when \mathbf{X} lives in a low-dimensional manifold, and (2) a sizable fraction of this information is used up in resolving minute perceptual details. These two points are closely related. First, the largeness of $S_{\text{NN}}^{\mathbf{X}} \approx D_{\text{KL}}(p_{\text{d}} \| p_{\text{eq}})$ is related to the fact that p_{eq} is supported on the entire volume, whereas p_{d} is restricted to the manifold, call it $\mathcal{M}_{\mathbf{X}}$. We can drill down further and triangulate the problem to the $\text{TC}(\mathbf{X})$ from Eq. (8). The joint density $p_{\text{d}}(x_1, \dots, x_{D_{\mathbf{X}}})$ is supported on $\mathcal{M}_{\mathbf{X}}$ whereas the product density $\prod_{k=1}^{D_{\mathbf{X}}} p_{\text{d}}(x_k)$ is supported over a larger joint space. Therefore, the KL between them diverges.

The same phenomenon can also occur in $I(X_1; X_2) = D_{\text{KL}}(p(x_1, x_2) \| p(x_1)p(x_2))$, where X_1 and X_2 are two 1-D continuous random variables. Unlike the discrete case, $I(X_1; X_2)$ can be infinite if X_1 and X_2 are perfectly correlated—specifying a real number requires *infinitely many digits of precision*, so we gain an infinite amount of information about X_1 from a given $X_2 = x_2$. More formally, the joint density $p(x_1, x_2)$ is confined to a lower-dimensional manifold, since $X_2 = f(X_1)$, whereas the product density $p(x_1)p(x_2)$ is supported over the full joint space (see App. A.1).

Extending the same intuition to $\text{TC}(\mathbf{X})$, we see that if the value of a pixel x_k resolves the intensity of any other pixel to very high precision, $\text{TC}(\mathbf{X})$ diverges. But this is what textures are: nearby pixels that are strongly correlated to one another. Under forward diffusion, these correlations are the first to vanish, since the less significant digits of every x_k are overwhelmed faster. In simpler terms, with a little noise added, we can still make out the semantic details in the image, but the smaller-scale details would have washed away (see App. F.3). It follows from this argument that $S_{\text{NN}}^{\mathbf{X}}$ should receive a large contribution from the neighborhood of $s = 0$ (or $t = T$). This is exactly what we observe empirically (see Fig. 14). The reverse drift term, and in particular $\nabla \log p(\mathbf{x}, t)$, becomes very large to confine the probability mass to the image manifold $\mathcal{M}_{\mathbf{X}}$.

6 EXPERIMENTS

Joint Gaussian Before delving into image models, it is worth considering a toy problem which allows precise control over $I(\mathbf{X}; \mathbf{Y})$ vis-à-vis $S_{\text{NN}}^{\mathbf{X}}$. Consider the random variables $\mathbf{X} \sim \mathcal{N}(0, \Sigma_{\mathbf{X}})$, and $\mathbf{Y} = \mathbf{A}\mathbf{X} + \varepsilon$ (see App. E.1). From Eqs. (55) and (56), we have the closed form expressions:

$$S_{\text{tot}}^{\mathbf{X}} \approx D_{\text{KL}}(p_{\text{d}} \| p_{\text{eq}}) = \frac{1}{2} (\text{Tr}(\Sigma_{\mathbf{X}}) - \log |\Sigma_{\mathbf{X}}| - D_{\mathbf{X}}), \quad I(\mathbf{X}; \mathbf{Y}) = \frac{1}{2} \log \left(\frac{|\Sigma_{\mathbf{Y}}|}{|\Sigma_{\varepsilon}|} \right), \quad (17)$$

where we have assumed a VP forward process, for which $p_{\text{eq}} = \mathcal{N}(0, I)$. The scores for this distribution are also known analytically, which allows us to compare the learned e_{θ} vectors with their idealized versions e_{θ}^* (cf. Eq. (9)). In particular, we choose

$$\Sigma_{\mathbf{X}} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{\top} + \epsilon I, \quad \mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_k, \overbrace{\lambda_{\delta}, \dots, \lambda_{\delta}}^{D_{\mathbf{X}} - k}), \quad \Sigma_{\varepsilon} = \sigma_{\varepsilon}^2 I. \quad (18)$$

where $\lambda_{\bullet} > 0$, \mathbf{U} is an orthogonal matrix and a small $\epsilon > 0$ ensures stability. This simple model allows us to study how well diffusion models absorb different amounts of $S_{\text{tot}}^{\mathbf{X}}$ and $I(\mathbf{X}; \mathbf{Y})$. We do

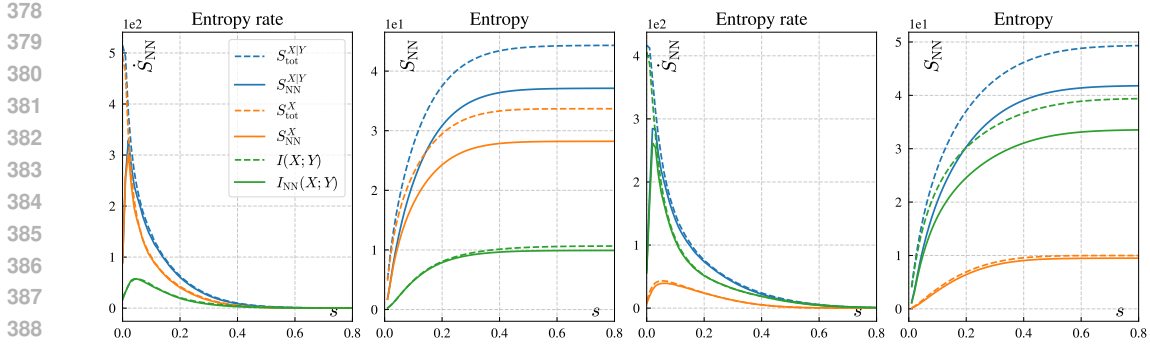


Figure 3: Two ways of having a large $S_{\text{tot}}^{\mathbf{X}|\mathbf{Y}}$. (Left) *Flattening*: $I(\mathbf{X}; \mathbf{Y})$ stays finite while $S_{\text{tot}}^{\mathbf{X}}$ blows up due to lower intrinsic dimensionality of \mathbf{X} . (Right) *Determinism*: $I(\mathbf{X}; \mathbf{Y})$ diverges when \mathbf{Y} is strongly correlated with \mathbf{X} , but $S_{\text{tot}}^{\mathbf{X}}$ is under control since its covariance is full-rank. The entropy rate curves are the time derivatives of the corresponding entropy or mutual information. Similar plots for different degrees of flattening and correlation are given in Figs. 9 and 10.

two kinds of experiments with this, corresponding to two different ways of making $S_{\text{tot}}^{\mathbf{X}|\mathbf{Y}}$ large. (1) *Flattening*: We adjust λ_δ to be much smaller than the other eigenvalues to induce an approximately low-rank structure in $\Sigma_{\mathbf{X}}$. This ‘flattens’ \mathbf{X} along $D_{\mathbf{X}} - k$ directions, which causes $S_{\text{tot}}^{\mathbf{X}}$ to blow up, simulating the manifold problem. At the same time we maintain $\Sigma_{\mathbf{Y}} = A\Sigma_{\mathbf{X}}A^T + \Sigma_\epsilon$ at full-rank by choosing Σ_ϵ appropriately, which keeps $I(\mathbf{X}; \mathbf{Y})$ under control (see Fig. 1). (2) *Determinism*: We keep $\Sigma_{\mathbf{X}}$ full rank and make \mathbf{Y} a more deterministic function of \mathbf{X} by making Σ_ϵ low (cf. Eq. (57)). Here, $S_{\text{tot}}^{\mathbf{X}}$ remains constant, whereas $I(\mathbf{X}; \mathbf{Y})$ diverges. The results are shown in Fig. 3.

We train two diffusion models, one on samples of \mathbf{X} alone, and the other on a joint sample of (\mathbf{X}, \mathbf{Y}) . In the flattening experiment, both models struggle to absorb the respective $S_{\text{tot}}^{\mathbf{X}}$ ’s, especially from early s , but they manage to estimate the $I(\mathbf{X}; \mathbf{Y})$ correctly! This is because the divergences in the conditional and unconditional scores cancel in Eq. (13), as evident from the green lines in the leftmost plot of Fig. 3. More intuitively, *the information required to locate the manifold $\mathcal{M}_{\mathbf{X}}$ is largely irrelevant in correlating \mathbf{X} with \mathbf{Y}* . On the other hand, if σ_ϵ is made small \mathbf{X} and \mathbf{Y} converge on the hyperplane $\mathbf{y} = A\mathbf{x}$. $\Sigma_{\mathbf{X}}$ is full rank, which keeps $S_{\text{tot}}^{\mathbf{X}}$ under control. The conditional and unconditional scores no longer cancel out, since the former grows large to confine the probability mass to the hyperplane in the joint $\mathbf{x}\mathbf{y}$ space, while the latter remains mild in comparison.

The flattening experiment is a simplified version of what happens in image models. The ratio of $I(\mathbf{X}; \mathbf{Y})$ to $S_{\text{tot}}^{\mathbf{X}}$ is much more extreme in the latter, hovering around 10^{-3} (see Fig. 14). Indeed, even in the Gaussian case, we observe that estimates of $I(\mathbf{X}; \mathbf{Y})$ deteriorate when λ_δ is made too small (see Fig. 11). Next, we study whether such coupling between $I(\mathbf{X}; \mathbf{Y})$ and the neural entropy can be detected in image models.

Images We do not have access to the true scores for image distributions, so we cannot gauge whether S_{NN} and $I(\mathbf{X}; \mathbf{Y})$ are close to their target values (see Fig. 14). Therefore, we take an alternate route to ascertain the claims in Sec. 5. Specifically, we want to assess whether the early time steps in s , those that pinpoint the manifold and resolve the minute details, contain any class-specific information. To do this, we allow the diffusion model to produce its own conditioning information by pairing it with an encoder $q_\phi(\mathbf{z}|\mathbf{x})$. The resulting arrangement, called a diffusion autoencoder (DAE), is discussed in great detail in App. F.1. We make some modifications.

Briefly, the DAE produces a compressed representation \mathbf{Z} of the images \mathbf{X} . The trick here is to note that we can condition a diffusion model on different \mathbf{Z} ’s at different stages of the diffusion process. We use two latents \mathbf{Z}_{per} and \mathbf{Z}_{sem} , and split the diffusion model loss into (see App. F.3 for details)

$$\mathcal{L}_{\text{DEM}}^{\text{split}} := \int_0^\tau ds \mathbb{E}_{\mathbf{Z}_{\text{per}}} [L(\mathbf{z}_{\text{per}}; s)] + \int_\tau^T ds \mathbb{E}_{\mathbf{Z}_{\text{sem}}} [L(\mathbf{z}_{\text{sem}}; s)], \quad (19)$$

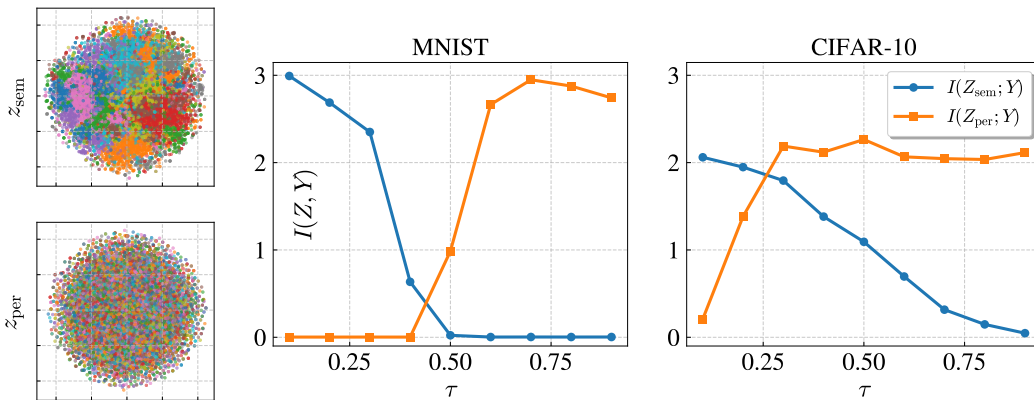


Figure 4: Probing the information stored in a diffusion model using a diffusion autoencoder. (Left column) t-SNE plots of z_{sem} and z_{per} for MNIST with $\tau = 0.1T$. The former shows discernible clusters corresponding to the different digits, while the latter has no such structure. This indicates that the information collected from $s \in (0, \tau)$ contains little to now information about the digits. (Right) The correlation between the learned latents and the true labels, as quantified by $I(\mathbf{Z}_\bullet; \mathbf{Y})$.

$$L(\mathbf{z}; s) := \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{x}}_s} \left[\frac{\sigma^2}{2} \left\| \nabla \log p_{\text{eq}}^{(s)}(\tilde{\mathbf{x}}_s) - \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) + \mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s; \mathbf{z}) \right\|^2 \right]. \quad (20)$$

The resulting arrangement encodes into \mathbf{Z}_{per} only the information that was absorbed from $s \in (0, \tau)$ by the diffusion model. The remaining information, from $s \in [\tau, T)$, imprints on \mathbf{Z}_{sem} . Thus, \mathbf{Z}_{per} and \mathbf{Z}_{sem} are compressed proxies for what the diffusion model ‘sees’ in those intervals. The DAE creates these latents from \mathbf{X} alone; it is not given the class labels \mathbf{Y} . Therefore, we can check how well \mathbf{Z}_{per} and \mathbf{Z}_{sem} correlate with \mathbf{Y} to understand which stages of diffusion are informed by \mathbf{Y} . This is done by estimating $I(\mathbf{Z}_{\text{per}}; \mathbf{Y})$ and $I(\mathbf{Z}_{\text{sem}}; \mathbf{Y})$ —the distributions of \mathbf{Z}_\bullet are not low-rank, so we train conditional diffusion models on $\mathbf{Z}_\bullet | \mathbf{Y}$ and use Eq. (13). Next, we scan τ from $0.1T \rightarrow 0.9T$, training a new DAE at each τ . The results are shown in Fig. 4, and Figs. 15 and 16.

A t-SNE diagram of the two latents immediately confirms our thesis: \mathbf{Z}_{per} shows little structure that indicates any class-specific clustering at $\tau = 0.1T$, indicating that there is little \mathbf{Y} -dependent information at early s . On the other hand, \mathbf{Z}_{sem} shows clear separation of clusters—the model can identify the digits from slightly noisy images of them. Furthermore, as we increase τ more \mathbf{Y} -specific information leaks into \mathbf{Z}_{per} from \mathbf{Z}_{sem} , as shown in the $I(\mathbf{Z}_\bullet, \mathbf{Y})$ plots. Thus, our insight from the flattening experiments also translates to image models: $I(\mathbf{X}; \mathbf{Y})$ is sourced from a different stage of the diffusion process than the one where the model strains to pin down the image manifold. This also explains why CFG is effective in image diffusion models; the CFG vector is strongest in the interval where the scores are well-behaved.

7 CONCLUSION

We have introduced a new perspective on diffusion models, interpreting them as a mechanism that learns correlations between the components of a signal \mathbf{X} itself, as well as between signals \mathbf{X} and \mathbf{Y} . To reproduce the former, the model must absorb $\text{TC}(\mathbf{X})$ worth of information, which is the non-trivial part of $S_{\text{tot}}^{\mathbf{X}}$. On top of this, an additional $I(\mathbf{X}; \mathbf{Y})$ information must be applied to correlate \mathbf{X} with \mathbf{Y} . When \mathbf{X} lives on a low-dimensional manifold, $\text{TC}(\mathbf{X})$ is high, but $I(\mathbf{X}; \mathbf{Y})$ is mostly decoupled from this divergence. This explains why CFG works, despite a low $I(\mathbf{X}; \mathbf{Y})$ to $S_{\text{tot}}^{\mathbf{X}}$ ratio.

Related work: Several works have studied mutual information in the context of diffusion models (Franzese et al. (2024); Kong et al. (2023; 2024); Wang et al. (2025)). However, this work is the first to study $\text{TC}(\mathbf{X})$ and the interplay between $\text{TC}(\mathbf{X})$ and $I(\mathbf{X}; \mathbf{Y})$. Furthermore, we attach a ‘physical’ meaning to these quantities, interpreting them as information *stored in the neural network*. This is an extension of the idea introduced in Premkumar (2025), where the control effort to transform $p_{\text{eq}} \rightarrow p_d$ is related to the information required to locate p_d in an ensemble centered at p_{eq} .

REFERENCES

- 486
487
488 Alexander A. Alemi, Ian Fischer, Joshua V. Dillon, and Kevin Murphy. Deep variational information
489 bottleneck. *CoRR*, abs/1612.00410, 2016. URL <http://arxiv.org/abs/1612.00410>.
490 16
- 491 Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and
492 Statistics. Springer, New York, 2006. ISBN 978-0387310732. 18, 23, 28
493
- 494 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal
495 Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao
496 Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
497 22
- 498 Arwen Bradley and Preetum Nakkiran. Classifier-free guidance is a predictor-corrector. In *NeurIPS*
499 *Workshop on Score-Based Methods*, 2024. URL <https://arxiv.org/abs/2408.09000>.
500 21
- 501 Kevin Clark and Priyank Jaini. Text-to-Image Diffusion Models are Zero-Shot Clas-
502 sifiers. In Alice Oh, Taesup Naumann, Amir Globerson, Kate Saenko, Moritz
503 Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing*
504 *Systems*, volume 36, pp. 58921–58937. Curran Associates, Inc., 2023. URL
505 [https://proceedings.neurips.cc/paper_files/paper/2023/file/
506 b87bdcf963cad3d0b265fcb78ae7d11e-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/b87bdcf963cad3d0b265fcb78ae7d11e-Paper-Conference.pdf). 6, 33
507
- 508 Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecom-*
509 *munications and Signal Processing)*. Wiley-Interscience, USA, 2006. ISBN 0471241954. 2,
510 16
- 511 Shaurya Dewan, Rushikesh Zavar, Prakanshul Saxena, Yingshan Chang, Andrew F. Luo,
512 and Yonatan Bisk. Diffusion PID: interpreting diffusion via partial information decom-
513 position. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ul-
514 rich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural In-*
515 *formation Processing Systems 38: Annual Conference on Neural Information Pro-*
516 *cessing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15,*
517 *2024*, 2024. URL [http://papers.nips.cc/paper_files/paper/2024/hash/
518 03d113a060c0ac93a5859517a0f07271-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2024/hash/03d113a060c0ac93a5859517a0f07271-Abstract-Conference.html). 21
- 519 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis.
520 In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.),
521 *Advances in Neural Information Processing Systems*, volume 34, pp. 8780–8794. Curran
522 Associates, Inc., 2021. URL [https://proceedings.neurips.cc/paper_files/
523 paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf). 1, 5
524
- 525 Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Associa-*
526 *tion*, 106(496):1602–1614, 2011. 18
- 527 Giulio Franzese, Mustapha Bounoua, and Pietro Michiardi. MINDE: Mutual in-
528 formation neural diffusion estimation. In *Proceedings of the International Con-*
529 *ference on Learning Representations (ICLR)*, pp. 16685–16716, 2024. URL
530 [https://proceedings.iclr.cc/paper_files/paper/2024/file/
531 47f75e809409709c6d226ab5ca0c9703-Paper-Conference.pdf](https://proceedings.iclr.cc/paper_files/paper/2024/file/47f75e809409709c6d226ab5ca0c9703-Paper-Conference.pdf). 2, 5, 9,
532 17, 20
- 533 Hao Ge and Da-Quan Jiang. Generalized jarzynski’s equality of inhomogeneous multidimen-
534 sional diffusion processes. *Journal of Statistical Physics*, 131(4):675–689, 5 2008. ISSN
535 1572-9613. doi: 10.1007/s10955-008-9520-4. URL [https://doi.org/10.1007/
536 s10955-008-9520-4](https://doi.org/10.1007/s10955-008-9520-4). 17
537
- 538 Shun-ichi Guo, Shlomo Shamai, S. Verdú, and J. J. Long. Mutual Information and Minimum Mean-
539 square Error in Gaussian Channels. *IEEE Transactions on Information Theory*, 51(4):1365–1381,
2005. 21

- 540 Irina Higgins, Loïc Matthey, Arka Pal, Christopher P. Burgess, Xavier Glorot, Matthew M.
541 Botvinick, Shakir Mohamed, and Alexander Lerchner. β -vae: Learning basic visual concepts
542 with a constrained variational framework. In *5th International Conference on Learning Repre-*
543 *sentations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. Open-
544 Review.net, 2017. URL <https://openreview.net/forum?id=Sy2fzU9gl>. 27
- 545 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *CoRR*, abs/2207.12598, 2022.
546 doi: 10.48550/ARXIV.2207.12598. URL [https://doi.org/10.48550/arXiv.2207.](https://doi.org/10.48550/arXiv.2207.12598)
547 [12598](https://doi.org/10.48550/arXiv.2207.12598). 2, 5, 22
- 548 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In
549 Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-
550 Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Con-*
551 *ference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12,*
552 *2020, virtual, 2020*. URL [https://proceedings.neurips.cc/paper/2020/hash/](https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html)
553 [4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html](https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html). 19
- 554 Chin-Wei Huang, Jae Hyun Lim, and Aaron C. Courville. A variational perspective on
555 diffusion-based generative models and score matching. In Marc’Aurelio Ranzato, Alina
556 Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Ad-*
557 *vances in Neural Information Processing Systems 34: Annual Conference on Neural Infor-*
558 *mation Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 22863–
559 22876, 2021. URL [https://proceedings.neurips.cc/paper/2021/hash/](https://proceedings.neurips.cc/paper/2021/hash/c11abfd29e4d9b4d4b566b01114d8486-Abstract.html)
560 [c11abfd29e4d9b4d4b566b01114d8486-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/c11abfd29e4d9b4d4b566b01114d8486-Abstract.html). 17, 20
- 561 David A. Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of*
562 *the IRE*, 40(9):1098–1101, 1952. doi: 10.1109/JRPROC.1952.273898. 19
- 563 Zahra Kadkhodaie, Florentin Guth, Eero P. Simoncelli, and Stéphane Mallat. Generalization in
564 diffusion models arises from geometry-adaptive harmonic representation. *CoRR*, abs/2310.02557,
565 2023. doi: 10.48550/ARXIV.2310.02557. URL [https://doi.org/10.48550/arXiv.](https://doi.org/10.48550/arXiv.2310.02557)
566 [2310.02557](https://doi.org/10.48550/arXiv.2310.02557). 27
- 567 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the Design Space
568 of Diffusion-Based Generative Models. In Sanmi Koyejo, S. Mohamed, A. Agar-
569 wal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Informa-*
570 *tion Processing Systems 35: Annual Conference on Neural Information Processing Sys-*
571 *tems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9,*
572 *2022, 2022*. URL [http://papers.nips.cc/paper_files/paper/2022/hash/](http://papers.nips.cc/paper_files/paper/2022/hash/a98846e9d9cc01cfb87eb694d946ce6b-Abstract-Conference.html)
573 [a98846e9d9cc01cfb87eb694d946ce6b-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/a98846e9d9cc01cfb87eb694d946ce6b-Abstract-Conference.html). 18
- 574 J. L. Kelly. A New Interpretation of Information Rate. *The Bell System Technical Journal*, 35(4):
575 917–926, 1956. doi: 10.1002/j.1538-7305.1956.tb03809.x. 1, 2
- 576 Diederik Kingma and Ruiqi Gao. Understanding Diffusion Objectives as the
577 ELBO with Simple Data Augmentation. In A. Oh, T. Naumann, A. Globerson,
578 K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Pro-*
579 *cessing Systems*, volume 36, pp. 65484–65516. Curran Associates, Inc., 2023. URL
580 [https://proceedings.neurips.cc/paper_files/paper/2023/file/](https://proceedings.neurips.cc/paper_files/paper/2023/file/ce79fbf9baef726645bc2337abb0ade2-Paper-Conference.pdf)
581 [ce79fbf9baef726645bc2337abb0ade2-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/ce79fbf9baef726645bc2337abb0ade2-Paper-Conference.pdf). 19
- 582 Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann
583 LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB,*
584 *Canada, April 14-16, 2014, Conference Track Proceedings, 2014*. URL [http://arxiv.org/](http://arxiv.org/abs/1312.6114)
585 [abs/1312.6114](http://arxiv.org/abs/1312.6114). 16, 17, 27
- 586 Diederik P. Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models.
587 *CoRR*, abs/2107.00630, 2021. URL <https://arxiv.org/abs/2107.00630>. 18
- 588 Xianghao Kong, Rob Brekelmans, and Greg Ver Steeg. Information-Theoretic Diffusion. In *The*
589 *Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda,*
590 *May 1-5, 2023*. OpenReview.net, 2023. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=UvmDCdSPDOW)
591 [UvmDCdSPDOW](https://openreview.net/forum?id=UvmDCdSPDOW). 9, 21

- 594 Xianghao Kong, Ollie Liu, Han Li, Dani Yogatama, and Greg Ver Steeg. Interpretable Diffusion
595 via Information Decomposition. In *Proceedings of the Twelfth International Conference on*
596 *Learning Representations (ICLR)*, 2024. URL <https://openreview.net/forum?id=X6tNkN6ate>. 2, 9, 21
- 598 Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical
599 report, University of Toronto, 2009. [https://www.cs.toronto.edu/~kriz/](https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf)
600 [learning-features-2009-TR.pdf](https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf). 22
- 602 Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied
603 to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.
604 726791. 22
- 605 Alexander C. Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your dif-
606 fusion model is secretly a zero-shot classifier. In *Proceedings of the IEEE/CVF International*
607 *Conference on Computer Vision (ICCV)*, pp. 2206–2217, October 2023. 6, 33
- 608 Marvin Li and Sitan Chen. Critical windows: Non-asymptotic theory for feature emergence in
609 diffusion models. In *Proceedings of the 41st International Conference on Machine Learning*
610 *(ICML)*, ICML’24, pp. 1097:1–1097:25. JMLR.org, 2024. 33
- 612 Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow
613 matching for generative modeling. In *The Eleventh International Conference on Learning*
614 *Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL
615 <https://openreview.net/forum?id=PqvMRDCJT9t>. 20
- 616 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios
617 of the data distribution. In *Forty-first International Conference on Machine Learning, ICML 2024,*
618 *Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL [https://openreview.](https://openreview.net/forum?id=CNicRIVIPA)
619 [net/forum?id=CNicRIVIPA](https://openreview.net/forum?id=CNicRIVIPA). 1
- 620 David J. C. MacKay. *Information Theory, Inference & Learning Algorithms*. Cambridge University
621 Press, USA, 2002. ISBN 0521642981. 2
- 623 Dimitra Maoutsa, Sebastian Reich, and Manfred Opper. Interacting particle solutions of fokker-
624 planck equations through gradient-log-density estimation. *Entropy*, 22(8):802, 2020. doi: 10.
625 3390/e22080802. URL <https://www.mdpi.com/1099-4300/22/8/802>. 24
- 626 Koichi Miyasawa. An empirical Bayes estimator of the mean of a normal population. *Bulletin of*
627 *the International Statistical Institute*, 38:181–188, 1961. 18
- 629 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin,
630 Ji-Rong Wen, and Chongxuan Li. Large Language Diffusion Models, 2025. URL <https://arxiv.org/abs/2502.09992>. 1
- 632 Michele Pavon. Stochastic control and nonequilibrium thermodynamical systems. *Applied Math-*
633 *ematics and Optimization*, 19(1):187–202, 1989. doi: 10.1007/BF01448198. URL <https://doi.org/10.1007/BF01448198>. 17
- 635 Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffu-
636 sion Autoencoders: Toward a Meaningful and Decodable Representation. In *Proceedings of the*
637 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10619–10629,
638 June 2022. 27, 28
- 640 Akhil Premkumar. Neural Entropy, 2025. URL <https://arxiv.org/abs/2409.03817>. 1,
641 4, 5, 9, 17
- 642 Galen Reeves, Henry D. Pfister, and Alex Dytso. Mutual information as a function of matrix SNR
643 for linear Gaussian channels. In *Proceedings of the 2018 IEEE International Symposium on*
644 *Information Theory (ISIT)*, pp. 1754–1758, 2018. doi: 10.1109/ISIT.2018.8437326. 15
- 645 R. Tyrrell Rockafellar and Roger J.-B. Wets. *Variational analysis / R. Tyrrell Rockafellar, Roger J.-*
646 *B. Wets*. Grundlehren der mathematischen Wissenschaften, 317. Springer, Berlin ;, 1998. ISBN
647 3540627723. URL <http://swbplus.bsz-bw.de/bsz063165805cov.htm>. 18

- 648 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
649 image segmentation. *CoRR*, abs/1505.04597, 2015. URL [http://arxiv.org/abs/
650 1505.04597](http://arxiv.org/abs/1505.04597). 20
- 651
- 652 D. L. Ruderman. The statistics of natural images. *Network: Computation in Neural Systems*, 5(4):
653 517–548, November 1994. doi: 10.1088/0954-898X/5/4/006. URL [https://dx.doi.org/
654 10.1088/0954-898X/5/4/006](https://dx.doi.org/10.1088/0954-898X/5/4/006). 30
- 655 J. J. Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. Cambridge University Press, Cam-
656 bridge, 3 edition, 2020. 19
- 657
- 658 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *CoRR*,
659 abs/2202.00512, 2022. URL <https://arxiv.org/abs/2202.00512>. 22
- 660
- 661 Udo Seifert. Entropy production along a stochastic trajectory and an integral fluctuation theorem.
662 *Phys. Rev. Lett.*, 95:040602, Jul 2005. doi: 10.1103/PhysRevLett.95.040602. URL [https://
663 link.aps.org/doi/10.1103/PhysRevLett.95.040602](https://link.aps.org/doi/10.1103/PhysRevLett.95.040602). 20
- 664 Udo Seifert. Stochastic Thermodynamics, Fluctuation Theorems and Molecular Machines. *Reports
665 on Progress in Physics*, 75(12):126001, Nov 2012. doi: 10.1088/0034-4885/75/12/126001. URL
666 <https://dx.doi.org/10.1088/0034-4885/75/12/126001>. 3
- 667
- 668 C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27
669 (3):379–423, 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x. 1
- 670 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
671 learning using nonequilibrium thermodynamics. In Francis Bach and David Blei (eds.), *Pro-
672 ceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings
673 of Machine Learning Research*, pp. 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. URL
674 <https://proceedings.mlr.press/v37/sohl-dickstein15.html>. 1, 22
- 675
- 676 Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum Likelihood
677 Training of Score-Based Diffusion Models. In Marc’Aurelio Ranzato, Alina Beygelz-
678 imer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances
679 in Neural Information Processing Systems 34: Annual Conference on Neural Informa-
680 tion Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 1415–
681 1428, 2021a. URL [https://proceedings.neurips.cc/paper/2021/hash/
682 0a9fdbb17feb6ccb7ec405cfb85222c4-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/0a9fdbb17feb6ccb7ec405cfb85222c4-Abstract.html). 17, 19
- 683 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and
684 Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In
685 *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria,
686 May 3-7, 2021*. OpenReview.net, 2021b. URL [https://openreview.net/forum?id=
687 PXTIG12RRHS](https://openreview.net/forum?id=PXTIG12RRHS). 1, 22
- 688 Edward O. Thorp. The Kelly Criterion in Blackjack Sports Betting, and the Stock Market. In
689 Leonard C MacLean, Edward O Thorp, and William T Ziemba (eds.), *THE KELLY CAPITAL
690 GROWTH INVESTMENT CRITERION THEORY and PRACTICE*, World Scientific Book Chap-
691 ters, chapter 54, pp. 789–832. World Scientific Publishing Co. Pte. Ltd., April 2011. URL
692 https://ideas.repec.org/h/wsi/wschap/9789814293501_0054.html. 2
- 693
- 694 Naftali Tishby, Fernando C. N. Pereira, and William Bialek. The Information Bottleneck Method.
695 *CoRR*, physics/0004057, 2000. URL <http://arxiv.org/abs/physics/0004057>. 16
- 696
- 697 S. Vaikuntanathan and C. Jarzynski. Dissipation and Lag in Irreversible Processes. *Europhysics
698 Letters*, 87(6):60005, oct 2009. doi: 10.1209/0295-5075/87/60005. URL [https://dx.doi.
699 org/10.1209/0295-5075/87/60005](https://dx.doi.org/10.1209/0295-5075/87/60005). 3
- 700 Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Ma-
701 chine Learning Research*, 9(86):2579–2605, 2008. URL [http://jmlr.org/papers/v9/
vandermaaten08a.html](http://jmlr.org/papers/v9/vandermaaten08a.html). 33

702 Chao Wang, Giulio Franzese, Alessandro Finamore, Massimo Gallo, and Pietro Michiardi. Infor-
703 mation theoretic text-to-image alignment. In *The Thirteenth International Conference on Learn-*
704 *ing Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL
705 <https://openreview.net/forum?id=Ugs2W5XFFo>. 9
706
707 Zhou Wang and Alan C. Bovik. Mean Squared Error: Love It or Leave It? A New Look at Signal
708 Fidelity Measures. *IEEE Signal Processing Magazine*, 26(1):98–117, 2009. doi: 10.1109/MSP.
709 2008.930649. 28
710
711 Satosi Watanabe. Information theoretical analysis of multivariate correlation. *IBM Journal of Re-*
712 *search and Development*, 4(1):66–82, 1960. doi: 10.1147/rd.41.0066. 2, 4
713
714 Yibo Yang, Justus C. Will, and Stephan Mandt. Progressive compression with universally quantized
715 diffusion models. In *The Thirteenth International Conference on Learning Representations, ICLR*
716 *2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=CxXGvKRDnL>. 19
717
718 Candi Zheng and Yuan Lan. Characteristic guidance: Non-linear correction for diffusion model at
719 large guidance scale. In *Forty-first International Conference on Machine Learning, ICML 2024,*
720 *Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=eOtjMYdGLt>. 21
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756 A INFORMATION THEORY

757 A.1 MUTUAL INFORMATION: A PRIMER

758 The mutual information between two random variables quantifies the reduction in uncertainty of one
759 variable given the value of the other variable. If X and Y are the random variables in question, then
760 the mutual information $I(X; Y)$ is the information gained about X through a measurement of Y . If
761 $p(x, y)$ is the joint distribution of X and Y and $p(x)$ and $p(y)$ are the marginals,
762

$$763 \begin{aligned} I(X; Y) &:= D_{\text{KL}}(p(x, y) \| p(x)p(y)) \\ &= S(X) - S(X|Y) = S(Y) - S(Y|X). \end{aligned} \quad (21)$$

764 By construction, $I(X; Y)$ is symmetric in its arguments, and it is non-negative. Mutual information
765 captures *all* forms of statistical dependence, not just linear ones. That said, it is easier to develop
766 some intuition for $I(X; Y)$ by considering a simple linear model

$$767 Y = aX + \varepsilon, \quad (22)$$

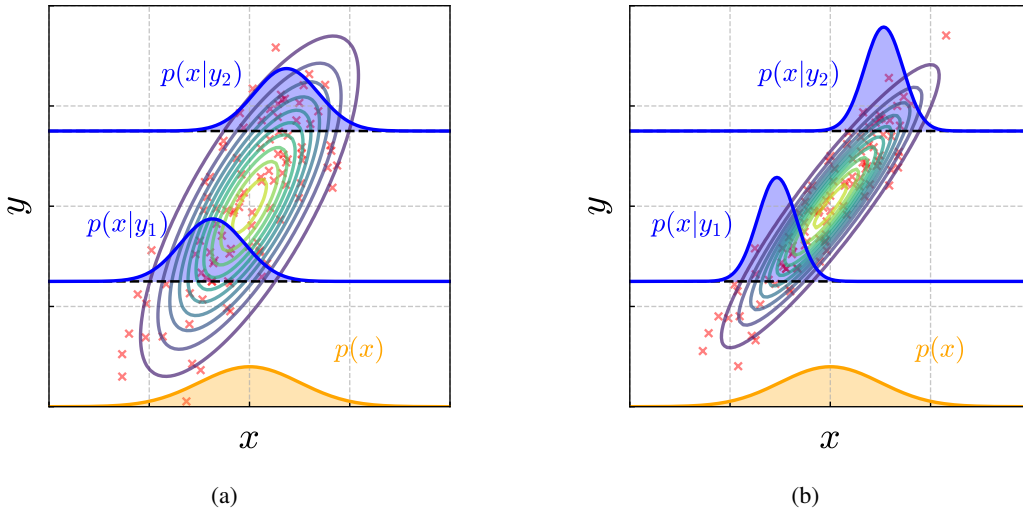
768 where $X \sim \mathcal{N}(0, \sigma_X^2)$, $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, and a is a real constant (Reeves et al., 2018). It is easy to see
769 that

$$770 Y|X = x \sim \mathcal{N}(ax, \sigma_\varepsilon^2), \quad (23a)$$

$$771 Y \sim \mathcal{N}(0, a^2\sigma_X^2 + \sigma_\varepsilon^2), \quad (23b)$$

$$772 X|Y = y \sim \mathcal{N}\left(\frac{a\sigma_X^2}{a^2\sigma_X^2 + \sigma_\varepsilon^2}y; \frac{\sigma_X^2\sigma_\varepsilon^2}{a^2\sigma_X^2 + \sigma_\varepsilon^2}\right), \quad (23c)$$

773 where Eq. (23a) follows from the fact that Y is a scaled version of X with some noise added to
774 it, and Eq. (23b) is obtained by marginalizing this distribution over X . With these distributions,
775 Eq. (23c) can be derived using Bayes' rule. Notice that $X|Y$ has a smaller variance than X . This
776 is what we mean when we say $p(x|y)$ is ‘narrower’ than $p(x)$ (see Fig. 5), although for general
777 distributions this is only true *on average*—there can be cases where the conditional is broader than
778 the marginal for some y .



801 Figure 5: The linear Gaussian model from Eq. (22) with (a) higher noise/larger σ_ε , and (b) lower
802 noise/smaller σ_ε . The blue curves are the conditionals $X|Y = y$ for some y , and the orange curve is
803 the marginal over X . Notice how the conditionals have a tighter variance compared to the marginal.
804 The contours are surfaces over constant probability in the joint distribution, and the red markers are
805 some samples.

806 The main point is that knowing Y dispels some of the uncertainty in X . That is, $X|Y = y$ has a
807 lower entropy on average than X ,

$$808 S(X|Y) = \int dy p(y) S(X|Y = y) = - \int dy p(y) \int dx p(x|y) \log p(x|y) \leq S(X). \quad (24)$$

Mutual information is the difference between these two entropies. We can compute the latter explicitly from Eqs. (23a) and (23b),

$$I(X; Y) = S(Y) - S(Y|X) = \frac{1}{2} \log \frac{\text{Var}(Y)}{\text{Var}(Y|X)} = \frac{1}{2} \log \left(1 + \frac{a^2 \sigma_X^2}{\sigma_\varepsilon^2} \right). \quad (25)$$

Notice that $I(X; Y) \rightarrow 0$ as $\sigma_\varepsilon \gg \sigma_X$, since the X signal is drowned out by the noise in this regime. In the opposite limit, when noise is very weak, X and Y are very strongly correlated and $I(X; Y)$ grows. If X and Y were discrete random variables, $I(X; Y)$ would have saturated at $S(X)$. However, in the continuous case, mutual information can diverge to infinity, which is the same pathology shared by differential entropy (Cover & Thomas, 2006, Chap. 8). Indeed, in the noiseless limit $p(x, y)$ collapses onto the line $y = ax$ whereas the product $p(x)p(y)$ spreads mass over the whole plane, so $p(x, y)$ is singular with respect to $p(x)p(y)$ in Eq. (21) (see Fig. 10).

This peculiar behavior of $I(X; Y)$ in the continuous case is reminiscent of the singular growth in entropy in image diffusion models as $t \rightarrow T$ (see Fig. 14). It is in fact the same phenomenon, which arises when the joint distribution converges on a lower dimensional manifold in the noiseless limit (see Fig. 6). In diffusion models the piece that becomes singular is the total correlation term in Eq. (8), $\text{TC}(\mathbf{X})$, which is a generalization of mutual information. We shall show in App. F.3 that the re-establishment of perceptual detail in the images coincides with a sharp peak of the neural entropy rate at the final stages of the generative process. Correlations between nearby pixels must be made very tight to get these small-scale details correct, which forces the image vector to track a manifold of lower dimensions, resulting in a divergent $\text{TC}(\mathbf{X})$. This is why the small details of the image take up a sizable portion of the total information budget.

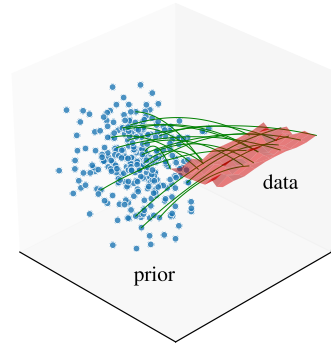


Figure 6

A.2 THE INFORMATION BOTTLENECK

Consider the problem of building a classifier that maps an input X to a label Y , where Y is now a discrete random variable. We would like to find a representation Z of X that captures all information in X that is relevant to predicting Y , while discarding the superfluous details. This is the viewpoint formalized by the information bottleneck, where the optimal assignment from X to Z is obtained by varying the stochastic map $q(z|x)$ to solve

$$\min_{q(z|x)} I(X; Z) - \gamma^{-1} I(Z; Y), \quad (26)$$

with $\gamma > 0$. A rigorous derivation of this functional is given in Tishby et al. (2000), but the intuition behind it is simple: minimizing $I(X; Z)$ maps a wide range of X values to a narrow range of the latent variable Z , and maximizing $I(Z; Y)$ creates a map between Z and Y where knowing Z almost completely determines Y . This forces $q(z|x)$ to only encode features from X into Z that most strongly associate with the class label Y , which is a form of selective compression. By adjusting γ we can set the tradeoff between compression and information preservation—setting $\gamma = 0$ collapses $q(z|x)$ to a single point, whereas $\gamma \rightarrow \infty$ pushes maximal detail from X to Z .

In general, Y can be any random variable, including a reconstruction of X itself. This is an autoencoder. However, there are a few subtle differences between Eq. (26) and the standard formulation of autoencoders (Kingma & Welling, 2014; Alemi et al., 2016). To see the connection, we expand Eq. (26) to

$$\min_{q(z|x)} S(Z) - S(Z|X) - \gamma^{-1} (S(Y) - S(Y|Z)). \quad (27)$$

The entropy of Y is independent of the encoder, so we can drop it from the objective (cf. Eq. (69)). We can introduce a regularization term to make the Z distribution close to a prior $p(z)$, like a standard normal distribution. This controls the entropy of Z from becoming too large as we vary $q(z|x)$. With these modifications Eq. (27) becomes

$$\min_{q(z|x)} \mathbb{E}_y \mathbb{E}_{q(z|x)} [-\log p(y|z)] + \gamma \mathbb{E}_x D_{\text{KL}}(q(z|x) \| p(z)). \quad (28)$$

Setting $Y = X$, we recover the negative ELBO from Eq. (66).

B STOCHASTIC CONTROL

In Sec. 3, we started with a formula for S_{tot} Eq. (4), which is derived in Premkumar (2025) using stochastic optimal control theory. In this section, we derive several known results about diffusion models following the same approach. This allows us to see how the Miyasawa/Tweedie relations, the denoising score matching objective, and the mutual information formula of Franzese et al. (2024), all spring from the same underlying optimization principle.

B.1 LOG LIKELIHOOD BOUND

In Sec. 3 we introduced the forward and reverse SDEs, Eqs. (2) and (3), which take $p_d \rightarrow p_0$ and back. If we replace the reverse process with an SDE (see Fig. 17)

$$d\mathbf{X}_t = -\mathbf{u}(\mathbf{X}_t, t)dt + \sigma(T-t)d\mathbf{B}_t, \quad (29)$$

and distribution $p_u(\cdot, 0)$ at time $t = 0$ evolves to $p_u(\cdot, T)$ at $t = T$, the log density of which is bound as (see Premkumar, 2025, App. E.2)

$$-\log p_u(\mathbf{x}, T) \leq \mathbb{E} \left[\left(\int_0^T ds \frac{\|b_+ - \mathbf{u}\|^2}{2\sigma^2} + \mathbf{u} \cdot \nabla \log p(\tilde{\mathbf{x}}_s, s | \tilde{\mathbf{x}}_0, 0) \right) - \log p_u(\tilde{\mathbf{x}}_T, 0) \Big| \tilde{\mathbf{X}}_0 = \mathbf{x} \right]. \quad (30)$$

The expectation value is taken over all trajectories generated by Eq. (2), starting at $\tilde{\mathbf{X}}_0 = \mathbf{x}$. This inequality, which can be derived from optimal control theory (Pavon, 1989), or the Feynman-Kac formula (Ge & Jiang, 2008; Huang et al., 2021), will be the starting point of many of our derivations. Completing the square, the bound can be written as

$$-\log p_u(\mathbf{x}, T) \leq \int_0^T ds \frac{1}{2\sigma^2} \mathbb{E} \left[\|b_+ - \sigma^2 \nabla \log p(\tilde{\mathbf{x}}_s, s | \tilde{\mathbf{x}}_0, 0) - \mathbf{u}\|^2 \Big| \tilde{\mathbf{X}}_0 = \mathbf{x} \right] \quad (31)$$

$$- \int_0^T ds \mathbb{E} \left[\frac{\sigma^2}{2} \|\nabla \log p(\tilde{\mathbf{x}}_s, s | \tilde{\mathbf{x}}_0, 0)\|^2 + \nabla \cdot b_+ \Big| \tilde{\mathbf{X}}_0 = \mathbf{x} \right] + S_0.$$

We have replaced $\mathbb{E}[-\log p_u(\tilde{\mathbf{x}}_T)]$ with the negative differential entropy $S_0 := \mathbb{E}_{p_0}[-\log p_0]$ by choosing $p_u(\cdot, 0)$ to be $p_0(\cdot)$ and noting that, to a very good approximation, $\tilde{\mathbf{x}}_T$ would be distributed as p_0 irrespective of the \mathbf{x} at which it started.

Averaging Eq. (31) over the data distribution p_d yields an upper bound on the cross-entropy between p_d and the reconstructed distribution $p_u(\cdot, T)$. In a diffusion model a neural network parametrizes the *control* \mathbf{u} , which affects only one term in the bound. Therefore, minimizing the cross-entropy is equivalent to minimizing the *denoising* objective

$$\mathbb{E}_{\mathbf{X}}[-\log p_u(\mathbf{x}, T)] + c \leq \int_0^T ds \frac{1}{2\sigma^2} \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{x}}_s} \left[\|b_+ - \sigma^2 \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) - \mathbf{u}\|^2 \right] := \mathcal{L}_D, \quad (32)$$

where c denotes the \mathbf{u} -independent terms from Eq. (31), averaged over \mathbf{X} . In an entropy-matching model $\mathbf{u} = -b_+ - \sigma^2 \mathbf{e}_\theta$, so the denoising entropy-matching objective is (cf. Eq. (67))

$$\mathbb{E}_{\mathbf{X}}[-\log p_\theta(\mathbf{x}, T)] + c \leq \int_0^T ds \frac{\sigma^2}{2} \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{x}}_s} \left[\left\| \nabla \log p_{\text{eq}}^{(s)}(\tilde{\mathbf{x}}_s) - \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) + \mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s) \right\|^2 \right] := \mathcal{L}_{\text{DEM}}. \quad (33)$$

A similar objective can be derived for score-matching models by setting $\mathbf{u} = b_+ - \sigma^2 \mathbf{s}_\theta$ in Eq. (32), or equivalently, $\mathbf{s}_\theta = \nabla \log p_{\text{eq}} + \mathbf{e}_\theta$ in Eq. (31) (Song et al., 2021a; Kingma & Welling, 2014).

B.2 OPTIMAL CONTROL AND REGRESSION

The bound in Eq. (30) is saturated by the *optimal control*,

$$\mathbf{u}_* = b_+ - \sigma^2 \nabla \log p, \quad (34)$$

which turns Eq. (29) into the reverse SDE Eq. (3), and the cross-entropy in Eq. (32) reaches its minimum value of $\mathbb{E}_{\mathbf{X}}[-\log p_d]$ (Pavon, 1989; Huang et al., 2021). But that also means \mathbf{u}_* minimizes the denoising objective \mathcal{L}_D ,

$$\mathbf{u}_* = \arg \min_{\mathbf{u}(\cdot, \cdot)} \mathcal{L}_D. \quad (35)$$

We apply Theorem 14.60 of Rockafellar & Wets (1998), which guarantees that the minimization of a time-integrated convex loss functional is achieved by pointwise minimization of the integrand. Briefly, given a normal integrand \mathcal{J} and a measurable weight function $\lambda(s) \geq 0$, the minimization of \mathcal{J} over the space χ of measurable functions $f : [0, T] \rightarrow \mathbb{R}^{D_{\mathbf{x}}}$ is

$$f_{\star} \in \arg \min_{f(\cdot) \in \chi} \int_0^T ds \lambda(s) \mathcal{J}(s, f(s)) = f_{\star}(s) \Leftrightarrow \arg \min_{f \in \mathbb{R}^{D_{\mathbf{x}}}} \mathcal{J}(s, f), \text{ for almost every } s \in [0, T]. \quad (36)$$

This allows us to analyze the denoising objective independently at each time s , which reduces Eq. (35) to a family of decoupled conditional mean regression problems, each minimizing the expected squared deviation at time s (see Sec. 1.5.5 of Bishop, 2006):

$$\begin{aligned} \mathbf{u}_{\star} &= \arg \min_{\mathbf{u} \in \mathbb{R}^{D_{\mathbf{x}}}} \int d\tilde{\mathbf{x}}_s \int d\mathbf{x} p(\tilde{\mathbf{x}}_s, \mathbf{x}) \|b_+(\tilde{\mathbf{x}}_s) - \sigma^2 \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) - \mathbf{u}(\tilde{\mathbf{x}}_s, s)\|^2 \\ &= b_+(\tilde{\mathbf{x}}_s) - \sigma^2 \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x} | \tilde{\mathbf{x}}_s)} \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0). \end{aligned} \quad (37)$$

Comparing with Eq. (34), we conclude that

$$\nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s) = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x} | \tilde{\mathbf{x}}_s)} \nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0). \quad (38)$$

If the perturbation kernel is Gaussian,³

$$\nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) = -\frac{\tilde{\mathbf{x}}_s - \mu(s)\mathbf{x}}{\Sigma(s)} \quad (39)$$

$$\stackrel{38}{\implies} \nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s) = -\frac{\tilde{\mathbf{x}}_s - \mu(s)\mathbb{E}[\mathbf{x} | \tilde{\mathbf{x}}_s]}{\Sigma(s)}, \quad (40)$$

which is the Miyasawa relation/Tweedie’s formula (Miyasawa, 1961; Efron, 2011). In simple terms, this relation states that at any given s , the ideal score is a vector pointing from $\tilde{\mathbf{x}}_s$ toward the denoised mean $\mathbb{E}[\mathbf{x} | \tilde{\mathbf{x}}_s]$, scaled by a forward factor. As mentioned above, Eq. (32) turns into an equality under Eq. (34), with $p_{\mathbf{u}_{\star}}(\cdot, T) = p_d(\cdot)$,

$$\mathbb{E}_{\mathbf{X}}[-\log p_d(\mathbf{x})] + c = \int_0^T ds \frac{\sigma^2}{2} \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{X}}_s} \left[\|\nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s) - \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0)\|^2 \right]. \quad (41)$$

For the Gaussian kernel, we can substitute Eqs. (39) and (40) and write this as (Kingma et al., 2021)

$$\begin{aligned} \mathbb{E}_{\mathbf{X}}[-\log p_d(\mathbf{x})] + c &= \int_0^T ds \frac{\sigma(s)^2}{2} \frac{\mu(s)^2}{\Sigma(s)^2} \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{X}}_s} \left[\|\mathbb{E}[\mathbf{x} | \tilde{\mathbf{x}}_s] - \mathbf{x}\|^2 \right] \\ &=: \int_0^T ds B(s) \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{X}}_s} \left[\|\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s) - \mathbf{x}\|^2 \right]. \end{aligned} \quad (42)$$

In the last step, we have collected the time-dependent prefactor into a single function $B(s)$, and defined $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s) := \mathbb{E}[\mathbf{x} | \tilde{\mathbf{x}}_s]$.

B.3 REWEIGHTED OBJECTIVE

Notice that the choice of weight function $\lambda(s)$ in Eq. (36) does not affect the pointwise minimization that leads to Eq. (38). This provides a theoretical justification of ‘variance dropping’ in practical

³ The kernel is Gaussian for Ornstein-Uhlenbeck processes, which have the form (Karras et al., 2022)

$$d\tilde{\mathbf{X}}_s = \phi(s)\tilde{\mathbf{X}}_s ds + \sigma(s)d\mathbf{B}_s.$$

The perturbation kernel of this SDE is

$$p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) = \mathcal{N}(\tilde{\mathbf{x}}_s; \mu(s)\mathbf{x}, \Sigma(s)I),$$

where

$$\mu(s) = \exp\left(\int_0^s d\bar{s} \phi(\bar{s})\right), \quad \Sigma(s) = \mu(s)^2 \int_0^s d\bar{s} \frac{\sigma(\bar{s})^2}{\mu(\bar{s})^2}.$$

denoising objectives such as Eq. (33) (Ho et al., 2020). That is, \mathcal{L}_{DEM} is replaced by the Monte Carlo average

$$T \mathbb{E}_{\mathbf{x} \sim p_d} \mathbb{E}_{s \sim \mathcal{U}(0, T)} \left[\lambda(s) \mathbb{E}_{\tilde{\mathbf{x}}_s \sim p(\tilde{\mathbf{x}}_s | \mathbf{x})} \left\| \nabla \log p_{\text{eq}}^{(s)}(\tilde{\mathbf{x}}_s) + \mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s) - \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) \right\|^2 \right], \quad (43)$$

where \mathcal{U} is the uniform distribution over $(0, T)$ and $\lambda(s)$ is not necessarily $\sigma^2(s)/2$. Dropping the variance means setting $\lambda(s) = 1$. In principle $\nabla \log p_{\text{eq}} + \mathbf{e}_\theta$ still recovers the optimal score from Eq. (38), but empirical observations show that alternative weighting schemes improve numerical stability and reduce gradient variance (Song et al., 2021a; Kingma & Gao, 2023). We used $\lambda(s) = 1$ in all our image diffusion models, including the DAE decoders. However, when applying Eq. (13) to estimate the mutual information we noticed that choosing $\lambda(s) = \sigma(s)^2/2$ gives slightly more accurate results.

C INFINITE VAEs AND QUANTUM MECHANICS

We can use Eq. (42) to develop a geometric view of diffusion models, which helps understand how they resolve the nuanced structure of complex distributions in small, continuous increments. For clarity, we focus on unconditional diffusion models for the moment. Recall that, for a Gaussian perturbation kernel,

$$\mathbb{E}_{\mathbf{X}}[-\log p_d(\mathbf{x})] + c = \int_0^T ds B(s) \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{X}}_s} \left[\|\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s) - \mathbf{x}\|^2 \right]. \quad (44)$$

Intuitively, $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s)$ is the average over all samples from p_d that had a *reasonable* probability of landing at $\tilde{\mathbf{x}}_s$ at time s under the forward SDE, Eq. (2). Each point in p_d can only travel so far under this process, so $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s)$ is the mean of points in p_d that are closest to $\tilde{\mathbf{x}}_s$, since those are the points that are most likely to arrive at $\tilde{\mathbf{x}}_s$ in finite time. Given $\tilde{\mathbf{x}}_s$, we can imagine firing off a swarm of stochastic trajectories from $\tilde{\mathbf{x}}_s$ back to $s = 0$, each of them evolving according to the reverse SDE, Eq. (3). The mean of the distribution of their endpoints is $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s)$ (see Fig. 18).

We can sharpen our intuition further by examining Eq. (44) for a single test point \mathbf{x} . That is, we lift the expectation over \mathbf{X} to obtain

$$-\log p_d(\mathbf{x}) + c'(\mathbf{x}) = \int_0^T ds B(s) \mathbb{E}_{\tilde{\mathbf{X}}_s} \left[\|\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s) - \mathbf{x}\|^2 | \tilde{\mathbf{X}}_0 = \mathbf{x} \right]. \quad (45)$$

Here c' is related to c as $c = \mathbb{E}_{\mathbf{X}}[c'(\mathbf{x})]$. Operationally, the r.h.s. can be evaluated through the following procedure: starting at \mathbf{x} , release a set of trajectories that follow the forward process, Eq. (2), and sample them at s to get a collection of $\tilde{\mathbf{x}}_s$ (see Fig. 18a). For each such $\tilde{\mathbf{x}}_s$, compute $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s)$, the mean of points in p_d that could have produced $\tilde{\mathbf{x}}_s$ under the forward process (see Figs. 18b and 18c). Therefore, the expectation over $\tilde{\mathbf{X}}_s$ is an average over the ℓ_2 -distance from the point \mathbf{x} to all candidates from p_d that are most similar to \mathbf{x} .

Propagating to $\tilde{\mathbf{x}}_s$ and traveling back to $\hat{\mathbf{x}}$ is how we locate candidates most like \mathbf{x} . The average in Eq. (45) admits a broader range of such options if s is larger. This is because the $\tilde{\mathbf{x}}_s$ samples at late s have little memory of where they started, so they could have come from almost anywhere in p_d as well. In other words, the integrand ‘sees more’ of p_d at larger s . In the very late time limit $\mathbb{E}[\mathbf{x} | \tilde{\mathbf{x}}_s] \approx \mathbb{E}[\mathbf{x}]$, the mean of p_d . Conversely, at smaller values of s we resolve the mean of a narrower range of the most \mathbf{x} -like points from p_d . Thus, Eq. (45) compares the test point with denoised means that become increasingly more specific as $s \rightarrow 0$ (see Fig. 19). Such gradual refinement echoes the logic of Huffman coding, in which symbols are distinguished by progressively finer splits (Huffman, 1952; Yang et al., 2025).

Connection to Quantum Mechanics At this point, readers with a physics background may notice an analogy to the *Fermi Golden Rule* for the transition rates between quantum states (Sakurai & Napolitano, 2020). With a mild abuse of the Dirac notation, the picture above can be formalized as

$$-\log p_d(\mathbf{x}) + c'(\mathbf{x}) = \int_0^T ds B(s) \int d\tilde{\mathbf{x}}_s \int d\mathbf{x}' \langle \mathbf{x}' | \mathcal{D}_s^\dagger | \tilde{\mathbf{x}}_s \rangle \langle \tilde{\mathbf{x}}_s | \mathcal{D}_s | \mathbf{x} \rangle \quad (46)$$

where \mathcal{D}_s is the operator that forward diffuses a delta function at \mathbf{x} to time s under Eq. (2), and \mathcal{D}_s^\dagger brings each $\tilde{\mathbf{x}}_s$ back to $s = 0$ according to Eq. (3). The average over \mathbf{x}' is the denoised mean $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s)$, toward which the score function at s points (cf. Eq. (40)). The integral over time and space aggregates all the $\mathbf{x} \rightarrow \mathbf{x}'$ transitions mediated by every possible $\tilde{\mathbf{x}}_s$.

In quantum mechanics, transition rates take a form similar to Eq. (46), with the \mathcal{D}_s operators replaced by the interaction term in the model. These rates are used to compute scattering cross-sections, which can be measured experimentally. This is how physicists test whether their theoretical model matches reality. If the predictions fail to agree with experiments, they must go back and construct a better model. The training of a diffusion model mirrors this procedure—given the training samples, which are observations from reality, we construct a model and iteratively refine it till it fits the data.

Writing Eq. (44) in the Dirac form, Eq. (46), allows us to interpret each $\mathbf{x} \rightarrow \mathbf{x}'$ transformation as an autoencoding step. That is, \mathcal{D}_s encodes $|\mathbf{x}\rangle$ into the intermediate states $|\tilde{\mathbf{x}}_s\rangle$, which can be viewed as a latent. Since forward diffusion is dissipative, $|\tilde{\mathbf{x}}_s\rangle$ contains less information than $|\mathbf{x}\rangle$, so \mathcal{D}_s is the first half of an information bottleneck (see App. A.2). On the other hand \mathcal{D}_s^\dagger decodes $|\tilde{\mathbf{x}}_s\rangle$ back to $|\mathbf{x}'\rangle$, the reconstructed version of $|\mathbf{x}\rangle$. At larger s the latent is less informative, and only the high-level details can be reconstructed. Thus, the integral over s implies that *a diffusion model is an infinite tower of variational autoencoders*, each capturing information at a different level of abstraction from the signal $|\mathbf{x}\rangle$. Shallower autoencoders in the tower, the ones at small s , capture the perceptual detail in the images, whereas the deeper ones retain the semantic features. We scan this tower with a DAE in App. F.

In stochastic thermodynamics, $-\log p_d(\mathbf{x})$ is interpreted as the sum total of the *path entropies* of each trajectory that starts from p_0 and ends at \mathbf{x} (Seifert, 2005). That is, it measures the accumulated information from all different ways of arriving at \mathbf{x} at time T . This is yet another interpretation of Eq. (46): given a test sample \mathbf{x} , we corrupt it by different degrees and aggregate the log probabilities of all paths that travel back to \mathbf{x} —it is proportional to the distance squared from the mean $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s)$ of reverse diffusion landings \mathbf{x}' . Thus, the geometric and entropic viewpoints converge.

It should be pointed out that our perspective differs from the one in Huang et al. (2021), where a diffusion model is viewed as a *single* infinitely deep autoencoder. They divide the time interval $(0, T]$ into infinitesimal steps, which are then interpreted as stochastic layers of this autoencoder. In contrast, Eq. (46) decomposes the diffusion model into a multitude of autoencoders, each one trying to reconstruct the signal at a different noise scale. This viewpoint aligns more closely with the simulation-free training of such models (Lipman et al., 2023). Furthermore, the autoencoders are not independent from one another—an encoder-decoder pair at scale s receives a subset of the information that flows through a pair at an earlier s . A harmonious synthesis of these two pictures is to view a diffusion model as the continuum limit of an autoencoder with *skip connections*, like a U-net (Ronneberger et al., 2015).

It should be stressed that Eq. (46) only holds for Gaussian kernels (cf. Eq. (40)), since a generic kernel does not admit the view that the score function points to the denoised mean (cf. Eqs. (38) and (41)). This can also be understood from the quantum perspective. The Fermi Golden rule is derived by keeping only the linear term in the perturbative Dyson series expansion. A Gaussian channel is also a linear perturbation, $\tilde{\mathbf{x}}_s = \mu(s)\mathbf{x} + \Sigma(s)\boldsymbol{\varepsilon}$ (cf. footnote 3), which is why the Fermi form Eq. (46) fully captures the transitions to intermediate states and back.

Finally, note that the above picture can be extended to conditional probabilities by simply replacing the marginal density with the conditional one everywhere. That is, $-\log p_d(\mathbf{x}|\mathbf{y})$ will be smaller if the denoised means $\hat{\mathbf{x}}(\tilde{\mathbf{x}}_s|\mathbf{y})$ are closer to \mathbf{x} , which is the case for conditional densities that assign high probability to regions close to \mathbf{x} .

D MUTUAL INFORMATION FROM DIFFUSION

D.1 MINDE

In Sec. 4 we discussed Eq. (13), a formula for mutual information originally derived in Franzese et al. (2024). We will give a derivation of this result using Eq. (30). Setting the control to its optimal

value, Eq. (34), and integrating by parts,

$$-\log p_d(\mathbf{x}) = \mathbb{E} \left[\left(\int_0^T ds \frac{\sigma^2}{2} \|\nabla \log p\|^2 - \nabla \cdot (b_+ - \sigma^2 \nabla \log p) \right) - \log p_0(\tilde{\mathbf{x}}_T) \Big| \tilde{\mathbf{X}}_0 = \mathbf{x} \right]. \quad (47)$$

Here $\nabla \log p \equiv \nabla \log p(\tilde{\mathbf{x}}_s, s)$. Averaging this over \mathbf{X} yields the entropy of $p_d(\mathbf{x})$ (cf. Eq. (31)),

$$S(\mathbf{X}) = \mathbb{E}_{\mathbf{X}}[-\log p_d(\mathbf{x})] = \int_0^T ds \mathbb{E}_{\tilde{\mathbf{x}}_s, \mathbf{X}} \left[\frac{\sigma^2}{2} \|\nabla \log p\|^2 - \nabla \cdot (b_+ - \sigma^2 \nabla \log p) \right] + S_0. \quad (48)$$

A similar formula can be derived for $S(\mathbf{X}|\mathbf{Y})$, by changing $\nabla \log p(\tilde{\mathbf{x}}_s, s) \rightarrow \nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y})$ and averaging over \mathbf{y} also. Mutual information is just the difference between the two,

$$\begin{aligned} I(\mathbf{X}; \mathbf{Y}) &= S(\mathbf{X}) - S(\mathbf{X}|\mathbf{Y}) \\ &= \mathbb{E}_{\mathbf{Y}} \left[\int_0^T ds \frac{\sigma^2}{2} \mathbb{E}_{\tilde{\mathbf{x}}_s, \mathbf{X}|\mathbf{y}} \left[\|\nabla \log p(\tilde{\mathbf{x}}_s, s)\|^2 - \|\nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y})\|^2 \right. \right. \\ &\quad \left. \left. + 2\nabla \cdot (\nabla \log p(\tilde{\mathbf{x}}_s, s) - \nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y})) \right] \right] \\ &\stackrel{\text{IBP}}{=} \mathbb{E}_{\mathbf{Y}} \left[\int_0^T ds \frac{\sigma^2}{2} \mathbb{E}_{\tilde{\mathbf{x}}_s, \mathbf{X}|\mathbf{y}} \left[\|\nabla \log p(\tilde{\mathbf{x}}_s, s)\|^2 - \|\nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y})\|^2 \right. \right. \\ &\quad \left. \left. - 2(\nabla \log p(\tilde{\mathbf{x}}_s, s) - \nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y})) \cdot \nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y}) \right] \right] \quad (49) \\ &= \mathbb{E}_{\mathbf{Y}} \left[\int_0^T ds \frac{\sigma^2}{2} \mathbb{E}_{\tilde{\mathbf{x}}_s, \mathbf{X}|\mathbf{y}} \left[\|\nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{y}) - \nabla \log p(\tilde{\mathbf{x}}_s, s)\|^2 \right] \right]. \quad (50) \end{aligned}$$

This is Eq. (13). We have assumed that T is sufficiently large that S_0 is nearly the same in both cases. We also partitioned the expectation over \mathbf{X} into an average over $\mathbf{X}|\mathbf{Y} = \mathbf{y}$ (shortened to $\mathbf{X}|\mathbf{y}$) and over \mathbf{Y} separately. This is why integration by parts in Eq. (49) produced the conditional score term.

It should be noted that Eq. (50) is a generalized version of Eq. (5) that appears in Kong et al. (2024), or Eq. (2) from Dewan et al. (2024), which holds only for Gaussian perturbation kernels. This relation is derived from the I-MMSE formula that appears in Guo et al. (2005); Kong et al. (2023), which relates the mutual information between variables \mathbf{X} and $\mathbf{Y} = \sqrt{\text{snr}}\mathbf{X} + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, I)$, to the minimum mean-squared error of estimating \mathbf{X} from \mathbf{Y} .

D.2 GUIDANCE

The expectation value in Eq. (48) is taken over $p(\tilde{\mathbf{x}}_s, s)$, which is the density forward evolved from the marginal $p_d(\mathbf{x})$. To compute the mutual information between \mathbf{Y} and the CFG-generated \mathbf{X} , we need to compute $S(\mathbf{X}|\mathbf{Y})_{\text{CFG}}$. The CFG modification from Eq. (16) is equivalent to replacing the score with

$$\nabla \log p(\mathbf{x}_t, t|\mathbf{y}) + w s_{\text{cl}}(\mathbf{y}, t) = \nabla \log \left(\frac{p(\mathbf{x}_t, t|\mathbf{y})^{1+w}}{p(\mathbf{x}_t, t)^w} \right) \quad (51)$$

in Eq. (3). But the resulting SDE is not the reversal of *anything* (Zheng & Lan, 2024; Bradley & Nakkiran, 2024). That is, there is no forward diffusion process for which the intermediate density is $p(\tilde{\mathbf{x}}_s, t|\mathbf{y})^{1+w}/p(\mathbf{x}_t, t)^w$. Samples of this distribution at $t = T$ are different from those of the true p_d (see Fig. 7). Therefore, we cannot write down an expression for $S(\mathbf{X}|\mathbf{Y})_{\text{CFG}}$ along the lines of Eq. (48), so we find no expression analogous to Eq. (50) for mutual information under CFG.

E EXPERIMENTS

We provide empirical evidence of the following claims made in the main text:

1. Conditional diffusion models store an additional amount of information, equal to the mutual information between \mathbf{X} and \mathbf{Y} (see Sec. 3 and Fig. 10). In nats,

$$I(\mathbf{X}; \mathbf{Y}) = S_{\text{tot}}^{\mathbf{X}|\mathbf{Y}} - S_{\text{tot}}^{\mathbf{X}} \approx S_{\text{NN}}^{\mathbf{X}|\mathbf{Y}} - S_{\text{NN}}^{\mathbf{X}}. \quad (52)$$

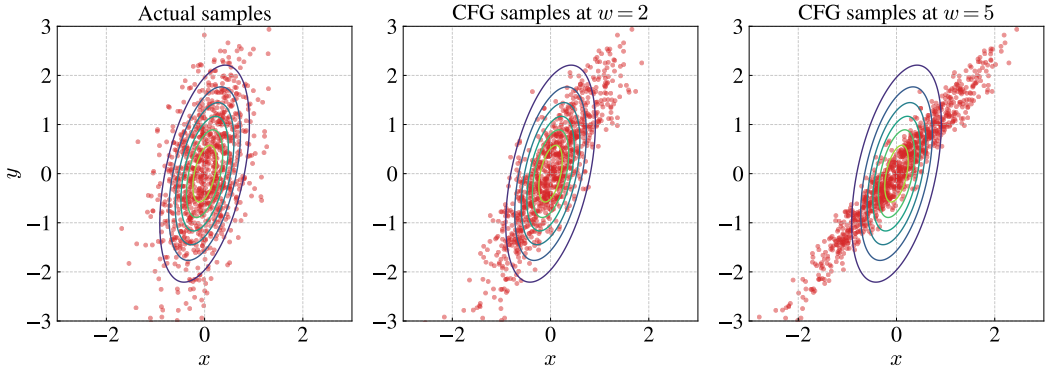


Figure 7: Samples generated by a CFG-style modification to the conditional score $\nabla \log p(x_t, t|y)$ of a joint Gaussian (cf. Eqs. (16) and (22)). CFG strengthens the correlation between X and Y , increasing their mutual information. But it also alters the relationship between them.

2. Neural entropy and $I(\mathbf{X}; \mathbf{Y})$ grow rapidly as \mathbf{X} and \mathbf{Y} become more strongly correlated, or when \mathbf{X} is low-dimensional. See discussions at the end of Sec. 4 and App. A.1 and Figs. 9 and 10.
3. CFG increases the mutual information between \mathbf{X} and \mathbf{Y} . See Sec. 4 and Fig. 8.
4. For images, the total information content is dominated by perceptual detail, which erodes rapidly in the first few forward diffusion steps. See App. F.2 and Figs. 13 and 14.
5. The perceptual information is largely the same for different image classes. Semantic structure is more closely correlated with the labels. See App. F.3 and Fig. 15.

The first three points can be demonstrated with a simple Gaussian model, like the one discussed in App. A.1. The mutual information and scores are known analytically, which allows us to compare the theoretical values of different entropies with their estimates from practical diffusion models. Image models are studied in App. F, by embedding them inside a DAE.

Our diffusion models used a U-net with self-attention layers (Ho & Salimans, 2022; Salimans & Ho, 2022), and were trained on H200 GPUs with 140 GB of memory. We used JAX/Flax as our ML framework (Bradbury et al., 2018), and trained our image models on the MNIST and CIFAR-10 datasets (LeCun et al., 1998; Krizhevsky, 2009). The Variance Preserving (VP) process was used in all experiments with diffusion models, for which $b_+(\tilde{x}_s, s) = -\beta(s)\tilde{x}_s/2$ and $\sigma(s) = \sqrt{\beta(s)}$ in Eq. (2) (Sohl-Dickstein et al., 2015; Song et al., 2021b). The code is available on GitHub.

E.1 A JOINT GAUSSIAN MODEL

We revisit the linear model from Eq. (22), generalizing it to higher dimensional random variables $\mathbf{X} \in \mathbb{R}^{D_X}$, $\mathbf{Y} \in \mathbb{R}^{D_Y}$. That is,

$$\mathbf{Y} = \mathbf{A}\mathbf{X} + \varepsilon, \tag{53}$$

where $\mathbf{X} \sim \mathcal{N}(0, \Sigma_X)$ and $\varepsilon \sim \mathcal{N}(0, \Sigma_\varepsilon)$. The joint vector $\mathbf{R} := (\mathbf{X}, \mathbf{Y})^\top$ is also Gaussian distributed, with zero mean, and covariance

$$\Sigma_R = \begin{pmatrix} \Sigma_X & \Sigma_X \mathbf{A}^\top \\ \mathbf{A} \Sigma_X & \mathbf{A} \Sigma_X \mathbf{A}^\top + \Sigma_\varepsilon \end{pmatrix} =: \begin{pmatrix} \Sigma_X & \Sigma_{\mathbf{X}\mathbf{Y}} \\ \Sigma_{\mathbf{X}\mathbf{Y}} & \Sigma_Y \end{pmatrix}. \tag{54}$$

This can be derived using $\Sigma_{\mathbf{X}\mathbf{Y}} \equiv \text{Cov}(\mathbf{X}, \mathbf{Y}) = \mathbb{E}[\mathbf{X}\mathbf{Y}^\top] - \mathbb{E}[\mathbf{X}]\mathbb{E}[\mathbf{Y}]^\top$ etc. The mutual information between \mathbf{X} and \mathbf{Y} is (cf. Eq. (25))

$$I(\mathbf{X}; \mathbf{Y}) = \frac{1}{2} \log \left(\frac{|\Sigma_Y|}{|\Sigma_{\mathbf{Y}|\mathbf{X}}|} \right), \tag{55}$$

where $\Sigma_{\mathbf{Y}|\mathbf{X}} = \Sigma_\varepsilon$ is just the average covariance of the distributions $\mathbf{Y}|\mathbf{X} \sim \mathcal{N}(\mathbf{A}\mathbf{x}, \Sigma_\varepsilon)$, and Σ_Y is defined in Eq. (54). If we use a VP process $p_{\text{eq}} = \mathcal{N}(0, I)$, and at sufficiently large T (cf.

Eq. (4),

$$S_{\text{tot}}^{\mathbf{X}} \approx D_{\text{KL}}(p_{\text{d}} \| p_{\text{eq}}) = D_{\text{KL}}(\mathcal{N}(0, \Sigma_{\mathbf{X}}) \| \mathcal{N}(0, I)) = \frac{1}{2} (\text{Tr}(\Sigma_{\mathbf{X}}) - \log |\Sigma_{\mathbf{X}}| - D_{\mathbf{X}}). \quad (56)$$

In one set of experiments we set $A \sim \mathcal{N}(0, 1)^{D_{\mathbf{X}} \times D_{\mathbf{Y}}}$, $\Sigma_{\epsilon} = \sigma_{\epsilon}^2 I$, and

$$\Sigma_{\mathbf{X}} = HH^{\top} + \epsilon I, \quad (57)$$

where $H \sim \mathcal{N}(0, 1)^{D_{\mathbf{X}} \times D_{\mathbf{X}}} / \sqrt{D_{\mathbf{X}}}$ and a small $\epsilon > 0$ ensures numerical stability as well as positive definiteness of $\Sigma_{\mathbf{X}}$. The $\Sigma_{\mathbf{X}}$ from Eq. (57) is full-rank, so $S_{\text{tot}}^{\mathbf{X}}$ is well-behaved. We vary σ_{ϵ} to adjust the correlation between \mathbf{X} and \mathbf{Y} , with $I(\mathbf{X}; \mathbf{Y})$ growing large as \mathbf{Y} becomes a more deterministic function of \mathbf{X} at small σ_{ϵ} (see Fig. 10).

In another set of experiments, we simulate the manifold problem in \mathbf{X} by setting

$$\Sigma_{\mathbf{X}} = U\Lambda U^{\top} + \epsilon I, \quad \Lambda = \text{diag}(\lambda_1, \dots, \lambda_k, \overbrace{\sigma_{\delta}, \dots, \sigma_{\delta}}^{D_{\mathbf{X}} - k}, \lambda_{\delta}), \quad (58)$$

where $\{\lambda_i\}_{i=1}^k \sim \mathcal{N}(0, 1) / \sqrt{k}$, and with smaller values of λ_{δ} effectively lowering the rank of \mathbf{X} . This time σ_{ϵ} is kept sufficiently large that $I(\mathbf{X}; \mathbf{Y})$ remains finite even as $S_{\text{tot}}^{\mathbf{X}}$ rapidly.

We can also compute the ideal score functions $\nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{y})$ and $\nabla \log p(\tilde{\mathbf{x}}_s, s)$ under the forward process. The conditional density at time s are obtained by evolving the initial distribution of $\mathbf{X} | \mathbf{Y} \sim \mathbf{y}$, namely (see Bishop, 2006, Sec. 2.3.1)

$$\mathcal{N}(\Sigma_{\mathbf{X}\mathbf{Y}}\Sigma_{\mathbf{Y}}^{-1}\mathbf{y}, \Sigma_{\mathbf{X}} - \Sigma_{\mathbf{X}\mathbf{Y}}\Sigma_{\mathbf{Y}}^{-1}\Sigma_{\mathbf{X}\mathbf{Y}}) =: \mathcal{N}(\mu_{\mathbf{X}|\mathbf{Y}}, \Sigma_{\mathbf{X}|\mathbf{Y}}). \quad (59)$$

We use the VP process in our experiments, under which

$$\tilde{\mathbf{x}}_s = \sqrt{\alpha(s)}\mathbf{x} + \sqrt{1 - \alpha(s)}\boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{N}(0, I). \quad (60)$$

Therefore, Eq. (59) is diffused to

$$p(\tilde{\mathbf{x}}_s, s | \mathbf{y}) = \mathcal{N}(\tilde{\mathbf{x}}_s; \mu_s^{\mathbf{X}|\mathbf{Y}}, \Sigma_s^{\mathbf{X}|\mathbf{Y}}), \quad (61)$$

$$\mu_s^{\mathbf{X}|\mathbf{Y}} := \mathbb{E}[\tilde{\mathbf{x}}_s | \mathbf{y}] = \sqrt{\alpha(s)}\Sigma_{\mathbf{X}\mathbf{Y}}\Sigma_{\mathbf{Y}}^{-1}\mathbf{y} \equiv \sqrt{\alpha(s)}\mu_{\mathbf{X}|\mathbf{Y}},$$

$$\begin{aligned} \Sigma_s^{\mathbf{X}|\mathbf{Y}} &:= \text{Cov}[\mathbf{X}_s | \mathbf{y}, \mathbf{X}_s | \mathbf{y}] \\ &= \mathbb{E}[(\mathbf{X}_s - \mu_s^{\mathbf{X}|\mathbf{Y}})(\mathbf{X}_s - \mu_s^{\mathbf{X}|\mathbf{Y}})^{\top}] \\ &= \alpha(s)\mathbb{E}[(\mathbf{X} - \mu_{\mathbf{X}|\mathbf{Y}})(\mathbf{X} - \mu_{\mathbf{X}|\mathbf{Y}})^{\top}] + (1 - \alpha(s))I \\ &= \alpha(s)\Sigma_{\mathbf{X}|\mathbf{Y}} + (1 - \alpha(s))I. \end{aligned}$$

Under Eq. (60), the marginal density at s is

$$p(\tilde{\mathbf{x}}_s, s) = \mathcal{N}(\tilde{\mathbf{x}}_s; 0, \Sigma_s^{\mathbf{X}}), \quad \Sigma_s^{\mathbf{X}} = \alpha(s)\Sigma_{\mathbf{X}} + (1 - \alpha(s))I. \quad (62)$$

Similarly, if the joint distribution were evolved by a VP process acting on both components of \mathbf{R} , the density at an intermediate time is

$$p(\tilde{\mathbf{x}}_s, \tilde{\mathbf{y}}_s, s) =: p(\tilde{\mathbf{r}}_s, s) = \mathcal{N}(\tilde{\mathbf{r}}_s; 0, \Sigma_s), \quad \Sigma_s^{\mathbf{R}} = \alpha(s)\Sigma_{\mathbf{R}} + (1 - \alpha(s))I. \quad (63)$$

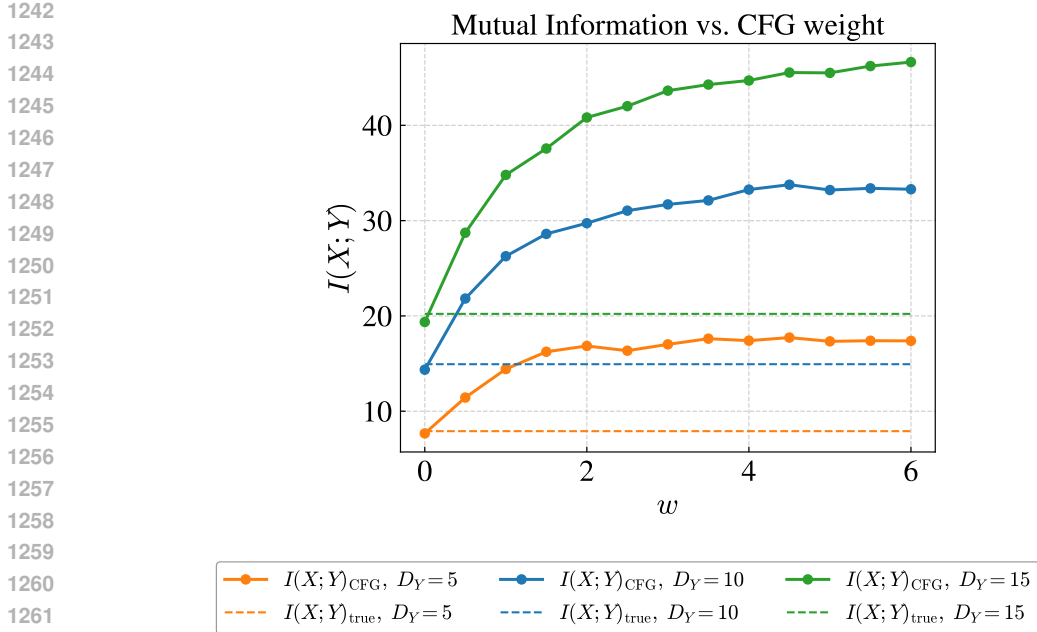
Then, the conditional, marginal, and joint scores are

$$\nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s | \mathbf{y}) = -(\Sigma_s^{\mathbf{X}|\mathbf{Y}})^{-1}(\tilde{\mathbf{x}}_s - \mu_s^{\mathbf{X}|\mathbf{Y}}), \quad (64a)$$

$$\nabla_{\tilde{\mathbf{x}}_s} \log p(\tilde{\mathbf{x}}_s, s) = -(\Sigma_s^{\mathbf{X}})^{-1}\tilde{\mathbf{x}}_s, \quad (64b)$$

$$\nabla_{\tilde{\mathbf{r}}_s} \log p(\tilde{\mathbf{r}}_s, s) = -(\Sigma_s^{\mathbf{R}})^{-1}\tilde{\mathbf{r}}_s. \quad (64c)$$

Equipped with these formulas, we can verify points 1 to 3 while sidestepping the singular behavior of neural entropy in image diffusion models. For each experiment with the joint Gaussian, a conditional diffusion model is trained on $\{(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\}_{i=1}^N$, and a second model learns the marginal of \mathbf{X} from $\{(\mathbf{x}^{(i)})\}_{i=1}^N$ alone. These models use a simple MLP core, and we train with the maximum likelihood denoising objective (Eq. (43) with $\lambda(s) = \sigma(s)^2/2$). The experiments are described below.



1263 Figure 8: Mutual information under CFG for the joint Gaussian model from Eq. (53). We fix
1264 $D_{\mathbf{X}} = 25$ and repeat the experiment with $D_{\mathbf{Y}} = 5, 10, 15$. Notice how $I(\mathbf{X}; \mathbf{Y})_{CFG}$ increases as
1265 the guidance strength is ramped up. It saturates faster for smaller $D_{\mathbf{Y}}$, when \mathbf{Y} has fewer degrees of
1266 freedom to encode the diversity in \mathbf{X} . This is also why the mutual information between images and
1267 labels is low in the first place. $I(\mathbf{X}; \mathbf{Y})_{CFG}$ was estimated using Eq. (14), with diffusion models
1268 trained on data generated with CFG. The true value of mutual information is known from Eq. (55).

1271 **Entropy and correlation** We train on samples of Eq. (53) for $D_{\mathbf{X}} = 25$, $D_{\mathbf{Y}} = 15$, with A kept
1272 fixed and $\sigma_{\varepsilon} = 1.0, 0.6, 0.25$. Reducing the noise strength increases $I(\mathbf{X}; \mathbf{Y})$ (cf. Fig. 5), as well
1273 as the conditional neural entropy $S_{NN}^{\mathbf{X}|\mathbf{Y}}$. We also plot the true value of these quantities calculated
1274 with the analytic scores in Eq. (64). The resulting entropy curves are shown in Fig. 10. Notice how
1275 the peak of the entropy rate curves becomes more localized at earlier s as the correlation between
1276 \mathbf{X} and \mathbf{Y} is made stronger. This is the same effect that gives rise to the sharp peak in the neural
1277 entropy rates for image diffusion models (see Fig. 14).

1278
1279 **Mutual information and guidance** In Sec. 4, we explained how CFG increases $I(\mathbf{X}; \mathbf{Y})$. For the
1280 joint Gaussian \mathbf{Y} is not a discrete random variable like a class label. Nonetheless, we can study how
1281 a ‘CFG-style’ modification to the reverse drift affects the samples from Eq. (53). Since we know the
1282 true scores, we can produce samples with the probability flow ODE (Maoutsa et al., 2020),

1283
1284
1285
1286

$$d\mathbf{x}_t = \left(-b_+(\mathbf{x}_t, T-t) + \frac{\sigma(T-t)^2}{2} [(1+w)\nabla \log p(\mathbf{x}_t, t|\mathbf{y}) - w\nabla \log p(\mathbf{x}_t, t)] \right) dt. \quad (65)$$

1287 A simple example of the samples generated by Eq. (65) is shown in Fig. 7. CFG tightens the
1288 dependence of \mathbf{X} on \mathbf{Y} , but also skews the true relationship between them (see App. D.2). Going to
1289 higher dimensions, we set $D_{\mathbf{X}} = 25$ and generate training data with Eq. (65) for $D_{\mathbf{Y}} = 5, 10, 25$,
1290 with a range of CFG weights $w \in (0, 6)$. We train a pair of diffusion models to reconstruct $\mathbf{X}|\mathbf{Y}$
1291 and \mathbf{X} , and estimate $I(\mathbf{X}; \mathbf{Y})$ using the MINDE formula, Eq. (14). The results are plotted in Fig. 8.

1292 As expected, CFG does increase the mutual information between \mathbf{X} and \mathbf{Y} . Two observations stand
1293 out: first, the increase in $I(\mathbf{X}; \mathbf{Y})$ saturates at larger w , and second, the gain in $I(\mathbf{X}; \mathbf{Y})$ is higher
1294 at larger $D_{\mathbf{Y}}$. Both these features can be understood through the information bottleneck principle
1295 from App. A.2: the degree to which the binding between \mathbf{X} and \mathbf{Y} can be strengthened is limited
by the number of degrees of freedom in \mathbf{Y} .

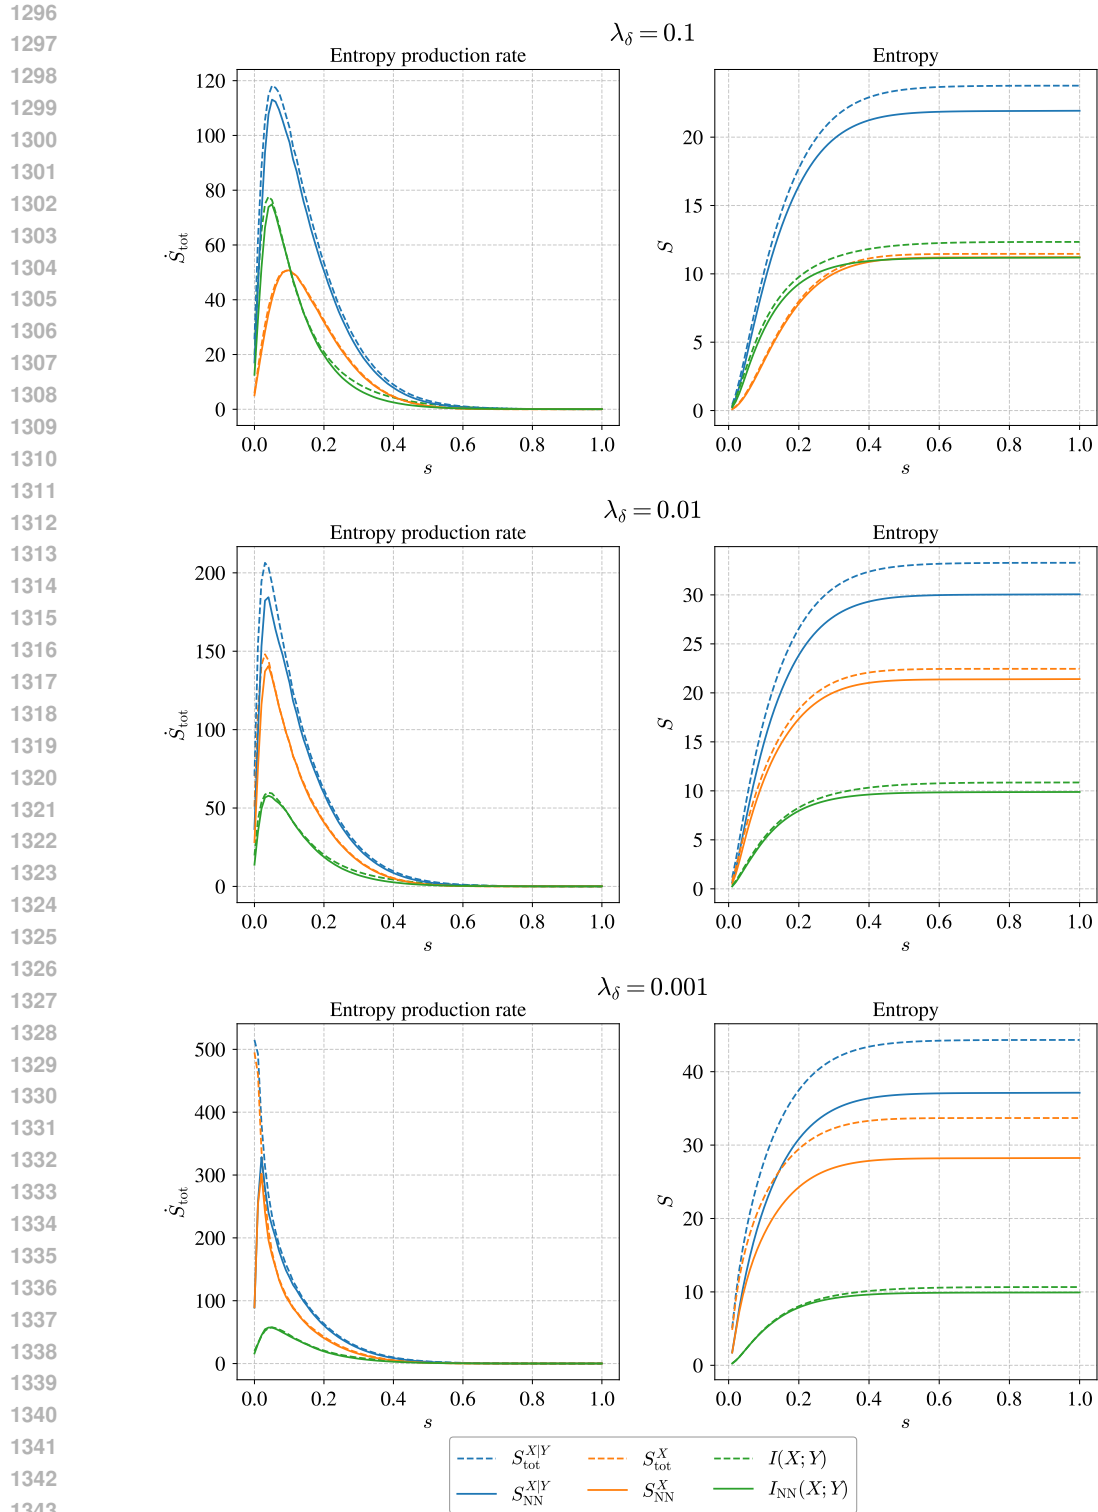


Figure 9: *Flattening*: Evolution of total entropy, neural entropy, and the mutual information under the forward process, for a joint Gaussian with $D_X = 25$, $D_Y = 15$, as the effective rank of Σ_X is made smaller. We keep $D_X = 25$, $D_Y = 15$, $\sigma_\epsilon = 1.0$ but λ_δ is lowered from top to bottom. This makes S_{tot}^X grow rapidly while $I(X; Y)$ remains finite. The bottom plot shows the diffusion model struggling to absorb information at the early time steps when S_{tot}^X and $S_{\text{tot}}^{X|Y}$ are both large. Despite this, the models compute $I(X; Y)$ more or less correctly. See Sec. 6 for discussion.

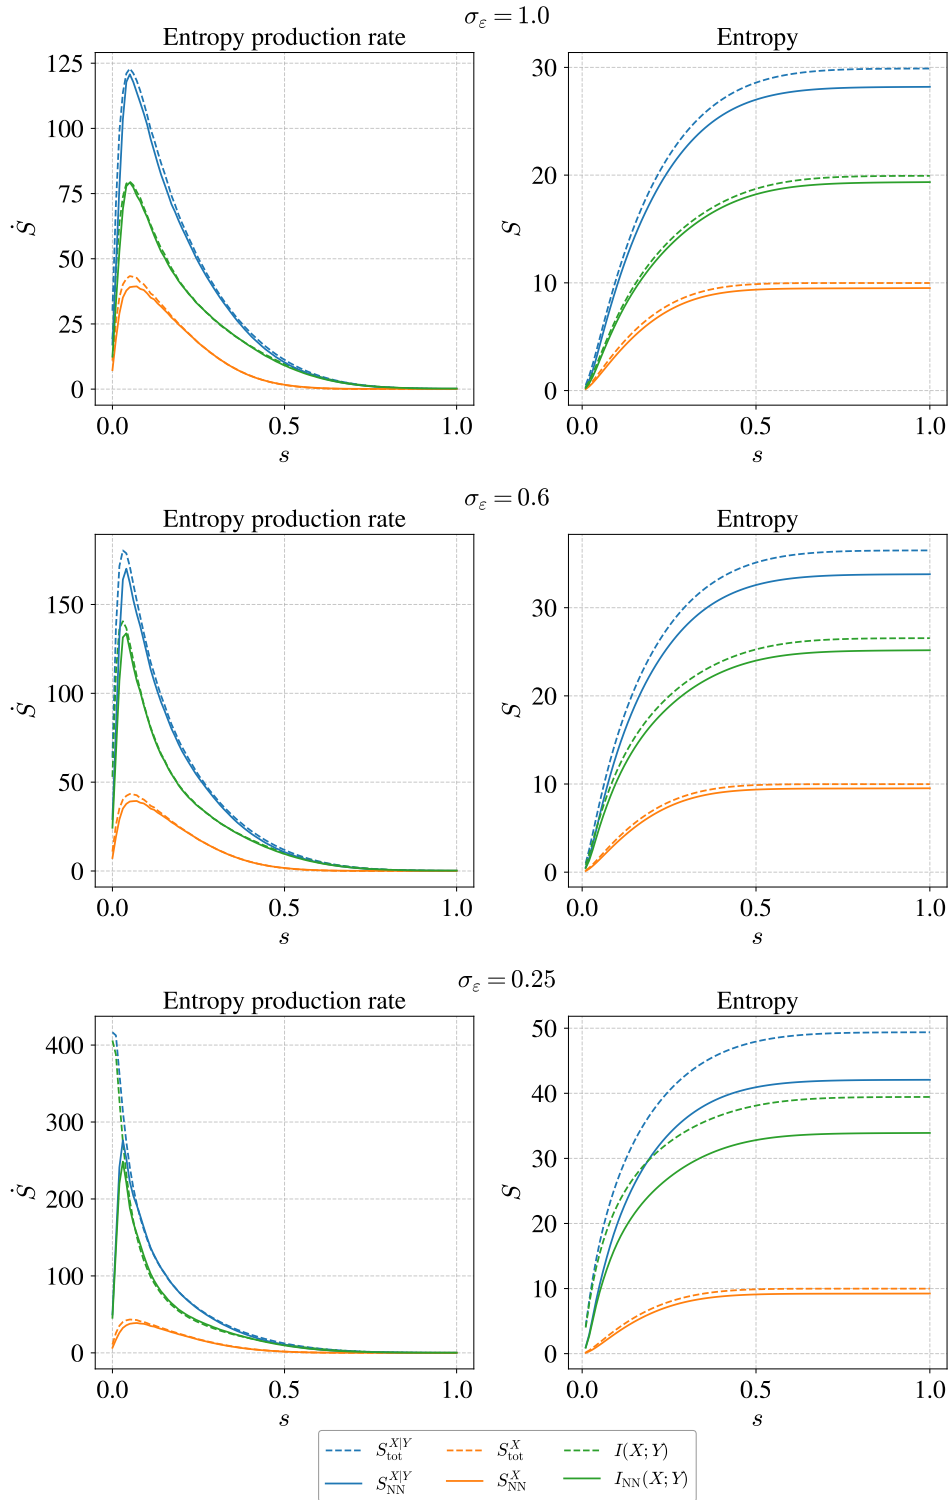


Figure 10: *Determinism*: Entropy curves for a joint Gaussian with $D_{\mathbf{X}} = 25$, $D_{\mathbf{Y}} = 15$ and a full-rank $\Sigma_{\mathbf{X}}$. As σ_{ϵ} is lowered \mathbf{X} and \mathbf{Y} become more correlated, which causes $I(\mathbf{X}; \mathbf{Y})$ to grow while S_{tot}^X remain fixed. Notice also how the $S_{\text{tot}}^{X|Y}$ and $I(\mathbf{X}; \mathbf{Y})$ curves become more concentrated near $s = 0$; as $\sigma_{\epsilon} \rightarrow 0$, \mathbf{X} and \mathbf{Y} converge on the hyperplane $\mathbf{y} = A\mathbf{x}$ which takes an infinite amount of information to locate precisely in the joint $\mathbf{x}\mathbf{y}$ space, even though \mathbf{X} itself is not lower-dimensional. The diffusion model no longer captures $I(\mathbf{X}; \mathbf{Y})$ accurately.

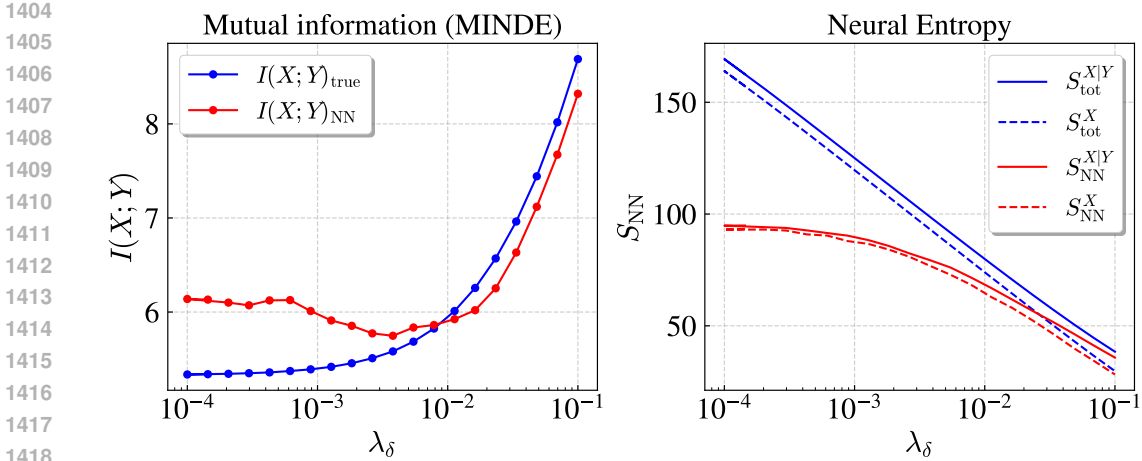


Figure 11: Mutual information and neural entropy for the joint Gaussian model as \mathbf{X} is flattened. We used $D_{\mathbf{X}} = 50, D_{\mathbf{Y}} = 8, k = 10$ and $\sigma_{\epsilon} = 1$. The blue lines are computed with the analytic expressions Eqs. (55) and (56). The variable λ_{δ} controls the spread of \mathbf{X} in $D_{\mathbf{X}} - k$ directions (cf. Eq. (58)). It is made smaller while \mathbf{Y} is maintained at full rank (see Fig. 9), which makes $S_{\text{tot}}^{\mathbf{X}}$ grow rapidly. The network struggles to keep pace at low λ_{δ} , as evident in the divergence between the S_{tot} and S_{NN} . In this range, even the $I(\mathbf{X}; \mathbf{Y})$ estimate begins to suffer.

F INFORMATION HIERARCHY

F.1 DIFFUSION AUTOENCODERS

A diffusion model can develop its own side information when it is paired with an encoder. This arrangement is called a *diffusion autoencoder*, or DAE (Preechakul et al., 2022). Recall that a standard variational autoencoder (VAE) is trained to minimize the negative of the evidence lower bound (Kingma & Welling, 2014),

$$-\text{ELBO}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[-\log p_{\theta}(\mathbf{x}|\mathbf{z})] + \gamma D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x}) \| p(\mathbf{z})) \quad (66)$$

where ϕ and θ are the encoder and decoder parameters respectively. The coefficient γ plays the role of the weighting factor in the β -VAE objective (Higgins et al., 2017), balancing reconstruction and KL terms. In a DAE, the reconstruction term is replaced by the upper bound

$$\mathbb{E}_{\mathbf{X}, \mathbf{Z}}[-\log p_{\theta}(\mathbf{x}|\mathbf{z})] + c \leq \int_0^T ds \mathbb{E}_{\mathbf{X}, \mathbf{Z}, \tilde{\mathbf{x}}_s} \left[\frac{\sigma^2}{2} \left\| \nabla \log p_{\text{eq}}^{(s)}(\tilde{\mathbf{x}}_s) - \nabla \log p(\tilde{\mathbf{x}}_s, s | \mathbf{x}, 0) + \mathbf{e}_{\theta}(\tilde{\mathbf{x}}_s, s; \mathbf{z}) \right\|^2 \right] =: \mathcal{L}_{\text{DEM}}^{\mathbf{X}|\mathbf{Z}}. \quad (67)$$

Here c is a constant with respect to the network parameters θ (cf. Eq. (33)). The expectation over \mathbf{X} and $\tilde{\mathbf{X}}_s$ averages over the data points $\{\mathbf{x}^{(i)}\}_{i=1}^N$ and their value at time s under the forward process in Eq. (2). Importantly, the bound is precisely the denoising entropy-matching objective used to train the diffusion model; minimizing this loss is equivalent to maximizing log likelihood (see App. B.1). A score-matching parameterization can also be used in Eq. (67). The latents \mathbf{z} are sampled from the encoder,

$$q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \mu_{\phi}(\mathbf{x}), \text{diag}(\sigma_{\phi}^2(\mathbf{x}))), \quad (68)$$

using the reparameterization trick to enable gradient-based training. Conditioning on \mathbf{z} allows the diffusion model to concentrate the probability mass to a smaller region in \mathbf{x} -space compared to the unconditional case; on average, conditional distributions are narrower than the marginals (cf. Eq. (11)). If the diffusion model had perfect freedom to choose the latent, it would assign a unique $\mathbf{z}^{(i)}$ to each $\mathbf{x}^{(i)}$ in the dataset, since that would lead to maximal concentration of probability in each conditional distribution. However, the DAE is unable to do so because (i) the inductive biases of the diffusion model temper its ability to perfectly resolve each $\mathbf{x}^{(i)}$, which is desirable because it avoids overfitting (Kadkhodaie et al., 2023), and (ii) the encoder admits a narrow range of \mathbf{z} , so the diffusion decoder has a limited set of latent codes to choose from—the DAE is an *information*

bottleneck (see App. A.2). Therefore, jointly minimizing the encoder term with the upper bound from Eq. (67) forces the diffusion model to negotiate a latent \mathbf{Z} that is maximally correlated with \mathbf{X} , under the given constraints. This follows from

$$\max I(\mathbf{X}; \mathbf{Z}) \equiv S(\mathbf{X}) - \min S(\mathbf{X}|\mathbf{Z}) \equiv S(\mathbf{X}) - \min \mathbb{E}_{\mathbf{X}, \mathbf{Z}}[-\log p_{\theta}(\mathbf{x}|\mathbf{z})], \quad (69)$$

since the cross entropy $\mathbb{E}_{\mathbf{X}, \mathbf{Z}}[-\log p_{\theta}(\mathbf{x}|\mathbf{z})]$ upper bounds the conditional entropy $S(\mathbf{X}|\mathbf{Z})$, and $S(\mathbf{X}) := \mathbb{E}_{\mathbf{X}}[-\log p_{\text{d}}(\mathbf{x})]$ is independent of θ or ϕ . The latent \mathbf{Z} is a compressed proxy for how the diffusion model represents \mathbf{X} . We use this fact in App. F, where the hierarchical nature of the information stored in these models is revealed through the structure they induce on the latents.

Minimizing Eq. (67) implicitly *maximizes* $S_{\text{NN}}^{\mathbf{X}|\mathbf{Z}}$, as evident from Eqs. (10) and (69)—a strongly correlated latent forces the diffusion model to discern tighter (on average) distributions of $\mathbf{X}|\mathbf{Z}$, which requires a higher neural entropy, whereas a weak latent does the opposite. This makes the DAE a great conceptual tool to understand how conditioning affects retention. Consider first the limiting case where the encoder is just the identity operator, so $\mathbf{Z} = \mathbf{X}$. There is now a unique $z^{(i)}$ for each $\mathbf{x}^{(i)}$, so every $p_{\theta}(\mathbf{x}|\mathbf{z})$ is a delta function, and the neural entropy $S_{\text{NN}}^{\mathbf{X}|\mathbf{Z}}$ is very large—the diffusion model has *memorized* each $\mathbf{x}^{(i)}$. At the other extreme, we can imagine an encoder that maps every value of $\mathbf{x}^{(i)}$ to a single value, call it z_{null} . This converts the decoder into an unconditional diffusion model since it receives no information about \mathbf{X} from the encoder. Consequently, the model learns to reconstruct the broadest possible distribution of \mathbf{X} , and $S_{\text{NN}}^{\mathbf{X}|\mathbf{Z}}$ reaches its lowest possible value, $S_{\text{NN}}^{\mathbf{X}}$. So the model retains the smallest amount of information when it is least committed to recovering each $\mathbf{x}^{(i)}$ perfectly.

This argument also connects to the tension between conditioning and generalization. In Sec. 4 we discussed the weak correlation between images \mathbf{X} and their labels \mathbf{Y} . If $I(\mathbf{X}; \mathbf{Y})$ were stronger, it would reduce the diversity in samples produced because the model has memorized more information. The power of CFG is that it is applied during the generative stage, so the model does not have to overcommit to the given data during training. However, CFG does have a fundamental limitation: if the underlying dependence between \mathbf{X} and \mathbf{Y} is weak, amplification of the signal can only go so far. DAE’s allow an alternative approach: a second diffusion model is trained to generate from \mathbf{Y} the latent \mathbf{Z} first, which is then used to produce \mathbf{X} . The latent \mathbf{Z} abstracts away the perceptual details that overwhelm the correlation between \mathbf{X} and \mathbf{Y} , while also being expressive enough to encode the variation in the semantic structure of \mathbf{X} (Preechakul et al., 2022).

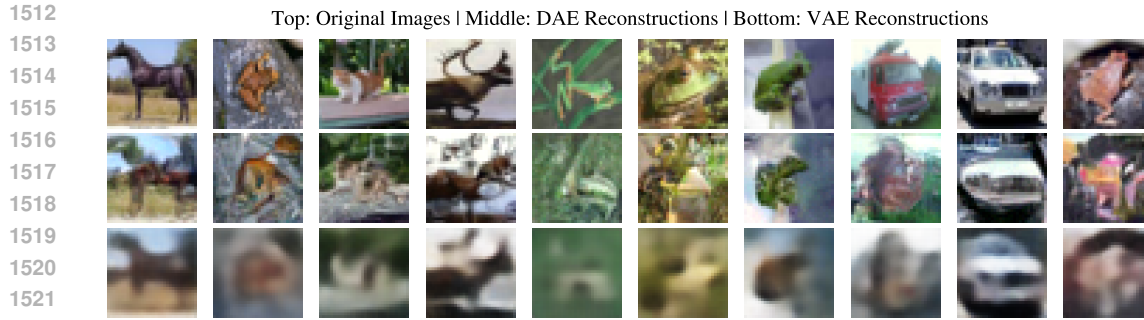
F.2 VAE vs. DAE

In App. F.1 we discussed diffusion autoencoders and pointed out that they help understand how the diffusion models store information. To see how this works, we start by comparing the diffusion model in the DAE with a simpler Gaussian-likelihood decoder, $p_{\psi}(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}; f_{\psi}(\mathbf{z}), \sigma_{\text{dec}}^2 I)$, where σ_{dec} is a constant and ψ are the network parameters. Minimizing the ℓ_2 loss of this decoder,

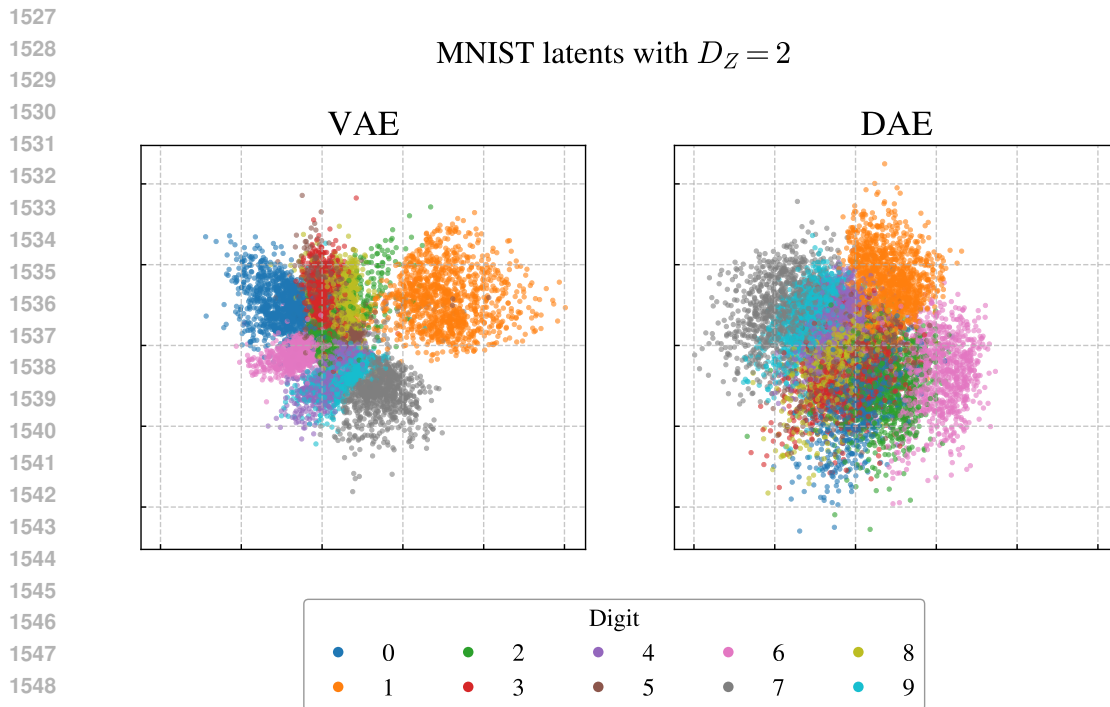
$$\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[-\log p_{\psi}(\mathbf{x}|\mathbf{z})] \propto \mathbb{E}_{\mathbf{z} \sim q_{\phi}(\cdot|\mathbf{x})}[\|\mathbf{x} - f_{\psi}(\mathbf{z})\|^2], \quad (70)$$

is equivalent to predicting the the conditional mean $\mathbb{E}[\mathbf{x}|\mathbf{z}]$, which lies between the modes of the true distribution (Bishop, 2006). As a result, in image processing applications, the reconstructions from such a decoder tend to be blurry (Wang & Bovik, 2009). On the other hand, the diffusion decoder from Eq. (67) generates a new sample by progressively evolving a random vector toward a denoised mean that becomes more resolved over time (see App. C and Fig. 19). Therefore, these models can capture the multi-modal structure of the underlying distribution with greater fidelity, producing reconstructions that are far more faithful to the original signal (see Fig. 12). Since the diffusion decoder retains more information about each \mathbf{x} , it can distinguish samples with greater accuracy. This places a greater strain on the encoder as it is pressured to supply more differentiated latent codes to disambiguate the richer variety of data points.

The latents in an autoencoder serve as a probe of the decoder’s ability to capture information. This is borne out in a simple experiment comparing the latents from a DAE to those from a VAE with the Gaussian decoder in Eq. (70), both of which use the same encoder architecture. We train both autoencoders to reconstruct MNIST images, restricting ourselves to latent dimensions of $D_{\mathbf{Z}} = 2$ for easier visualization of the aggregated posterior, $q_{\phi}(\mathbf{z}) = \sum_{\mathbf{x}} q_{\phi}(\mathbf{z}|\mathbf{x})p_{\text{d}}(\mathbf{x})$. Even at this low $D_{\mathbf{Z}}$, we observe discernible clustering in the VAE latent, corresponding to the different digits. By contrast, in the latent space of the DAE the clusters are more blended, with weaker separation



1523 Figure 12: Images reconstructed by a DAE and VAE. Both of them have the same encoder architecture.
1524 The VAE uses a Gaussian decoder that tends to produce blurrier outputs, whereas the diffusion
1525 decoder captures significantly more textural detail, leading to sharper images. In this example, the
1526 convolutional encoder’s simplicity limits the fidelity of the DAE reconstruction.



1550 Figure 13: Latents from a VAE and DAE trained to reconstruct MNIST digits. Distinct clusters
1551 appear in the VAE latent, even at the low dimensionality of $D_Z = 2$. On the other hand, the DAE’s
1552 latent clusters are more blended because the small-scale details captured by the diffusion decoder are
1553 similar for all digits, and this information overwhelms the semantic differences. See also Fig. 15a.

1554
1555
1556 between digit classes (see Fig. 13). This suggests that the DAE perceives greater similarity between
1557 different digits than the VAE, the common information across digit classes being the high-frequency
1558 detail washed away by the averaging effect of the Gaussian decoder. If we widen the bottleneck
1559 by increasing D_Z , we find better separation between the DAE clusters, since there is more room to
1560 encode the rich detail preserved by the diffusion decoder.

1561 The above experiment gives us a clue as to why image diffusion models often neglect conditioning
1562 on class labels. The *semantic* information that identifies the digit ‘1’ from an image of ‘1’ is a
1563 relatively small fraction of the total information content in that image. The rest encodes *perceptual*
1564 details that have a similar distribution for all images, even those of different digits. Therefore, the
1565 marginal \mathbf{X} and the conditional $\mathbf{X} | \mathbf{Y} = \mathbf{y}$ possess comparable entropy—specifying the class label
does not reduce the uncertainty in \mathbf{X} by a lot. In other words, the mutual information between these

images and their labels is low; the problem lies in the data itself. CFG is a trick to boost $I(\mathbf{X}; \mathbf{Y})$ post-training, but it merely amplifies whatever signal is already present; *multiplying a weak signal also magnifies the noise*.

F.3 SEMANTIC VS. PERCEPTUAL

Why must the diffusion model devote a large fraction of its information budget to resolving the microscopic details of the image? And how do we know it is these details that overwhelm the semantic information? To answer these questions, we begin by noting that forward diffusion dissolves the perceptual details in the first few steps, whereas the semantic structure is preserved—we can still read off a digit from a noisy image of it. More prosaically, natural images follow a power-law spectrum, which means the low frequencies dominate while high frequency (short wavelength) modes are subdued (Ruderman, 1994). Since the white noise term in Eq. (2) injects equal power across all frequencies, the finer details fade away more rapidly when images are diffused. Therefore, we expect entropy production associated with the removal of perceptual detail to be localized in a narrow interval near $s = 0$. The neural entropy rates in Fig. 14 exhibit a sharp peak in this range, which answers the second question.

We can also understand Fig. 14 from a geometric perspective by viewing S_{NN} as the information the network injects in the t -direction. The data distribution resides on a low-dimensional submanifold of the ambient pixel space. The reduced dimensionality of the data manifold stems from the fact that nearby pixels are very strongly correlated in high-fidelity images, so there are fewer degrees of freedom than the naive pixel count. In the generative stage, the diffusion model drives a high-dimensional Gaussian distribution back onto the lower-dimensional data manifold (see Fig. 6). The sharp rise in entropy rate as $t \rightarrow T$ is reflective of the fact that the network must supply substantial information to locate the manifold exactly, which involves collapsing the distribution to delta functions along all directions orthogonal to its tangent space. This singular behavior can also be traced back to the total correlation term in Eq. (8), as explained in Sec. 4, and at the end of App. A.1. This addresses the first question.

In App. C we conceptualized a diffusion model as an infinite tower of autoencoders, one for each instant s in the forward diffusion process. The encoder was denoted by the operator \mathfrak{D}_s and the decoder by \mathfrak{D}_s^\dagger . The shallow/small- s autoencoders are responsible for the small-scale details, whereas the deeper ones attend to the macroscopic features. It is possible to peer into this tower using a DAE, by conditioning its diffusion model on separate latents over different intervals in s . Recall Eq. (67), which we shall write as

$$\mathbb{E}_{\mathbf{X}, \mathbf{Z}}[-\log p_\theta(\mathbf{x}|\mathbf{z})] + c(T) \leq \int_0^T ds \mathbb{E}_{\mathbf{Z}}[L(\mathbf{z}; s)], \quad (71)$$

$$L(\mathbf{z}; s) := \mathbb{E}_{\mathbf{X}, \tilde{\mathbf{X}}_s} \left[\frac{\sigma^2}{2} \left\| \nabla \log p_{\text{eq}}^{(s)}(\tilde{\mathbf{x}}_s) - \nabla \log p(\tilde{\mathbf{x}}_s, s|\mathbf{x}, 0) + \mathbf{e}_\theta(\tilde{\mathbf{x}}_s, s; \mathbf{z}) \right\|^2 \right]. \quad (72)$$

Notice that if the integrate L from an intermediate time $s = \tau$ up to T the l.h.s. must be updated with the reconstructed density at τ ,

$$\mathbb{E}_{\tilde{\mathbf{X}}_\tau, \mathbf{z}_{\text{sem}}}[-\log p_\theta(\tilde{\mathbf{x}}_\tau, \tau|\mathbf{z}_{\text{sem}})] + c(\tau) \leq \int_\tau^T ds \mathbb{E}_{\mathbf{Z}}[L(\mathbf{z}_{\text{sem}}; s)], \quad (73)$$

where $c(\tau)$ is still independent of θ . The latent \mathbf{z}_{sem} encodes $\tilde{\mathbf{X}}_\tau$, the version of \mathbf{X} that has been forward diffused for a time τ . In other words, \mathbf{z}_{sem} represents the information stored in the $\mathfrak{D}_s/\mathfrak{D}_s^\dagger$ autoencoders for $s \geq \tau$. Following our earlier logic, \mathbf{z}_{sem} manages to evade much of the perceptual information—it ‘sees’ images where most of these microscopic details have been washed out and only the semantic structure remains—if τ is chosen judiciously. We can introduce another latent, \mathbf{z}_{per} , to aggregate the information from the $(0, \tau)$. Therefore, Eq. (71) can be split into

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}, \mathbf{z}_{\text{sem}}, \mathbf{z}_{\text{per}}}[-\log p_\theta(\mathbf{x}|\{\mathbf{z}_{\text{sem}}, \mathbf{z}_{\text{per}}\})] + c(T) \\ & \leq \int_0^\tau ds \mathbb{E}_{\mathbf{z}_{\text{per}}}[L(\mathbf{z}_{\text{per}}; s)] + \int_\tau^T ds \mathbb{E}_{\mathbf{z}_{\text{sem}}}[L(\mathbf{z}_{\text{sem}}; s)] =: \mathcal{L}_{\text{DEM}}^{\text{split}}. \end{aligned} \quad (74)$$

Thus, \mathbf{z}_{sem} and \mathbf{z}_{per} access information from different epochs of the forward diffusion process. We can verify points 4 and 5 by examining each of these latents closely.

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

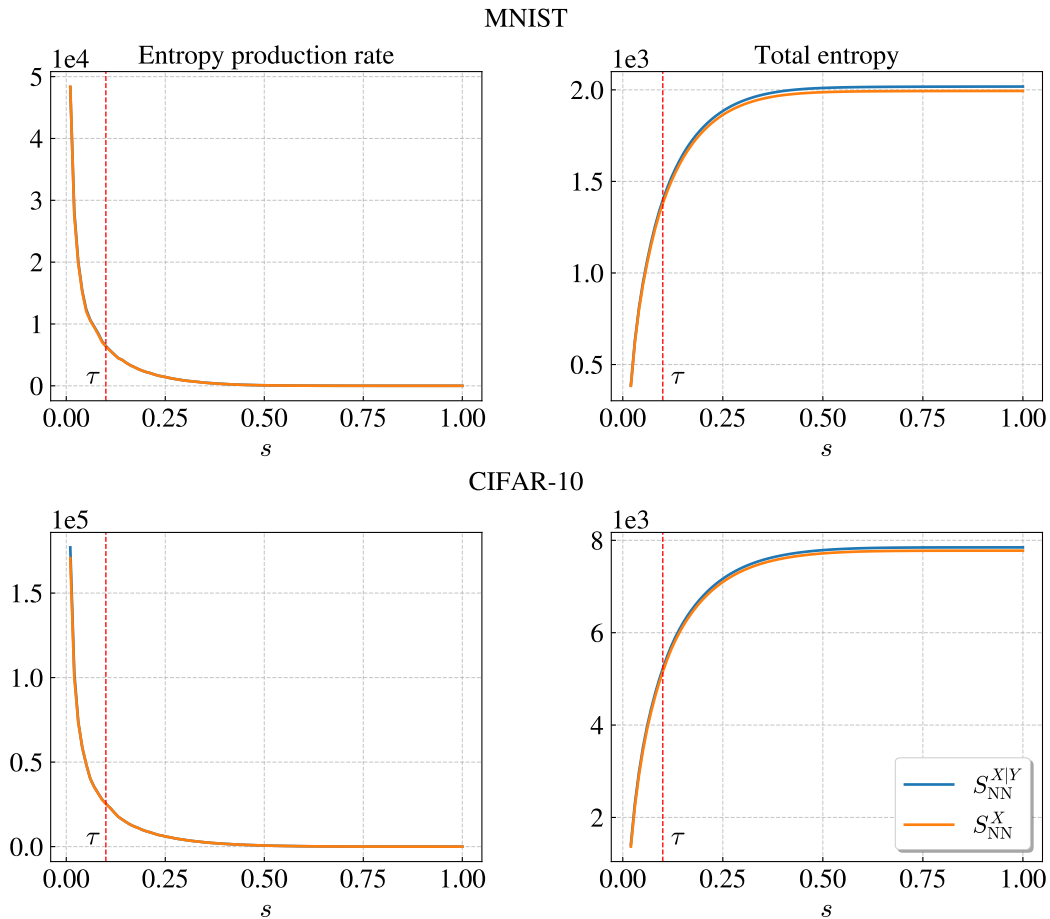
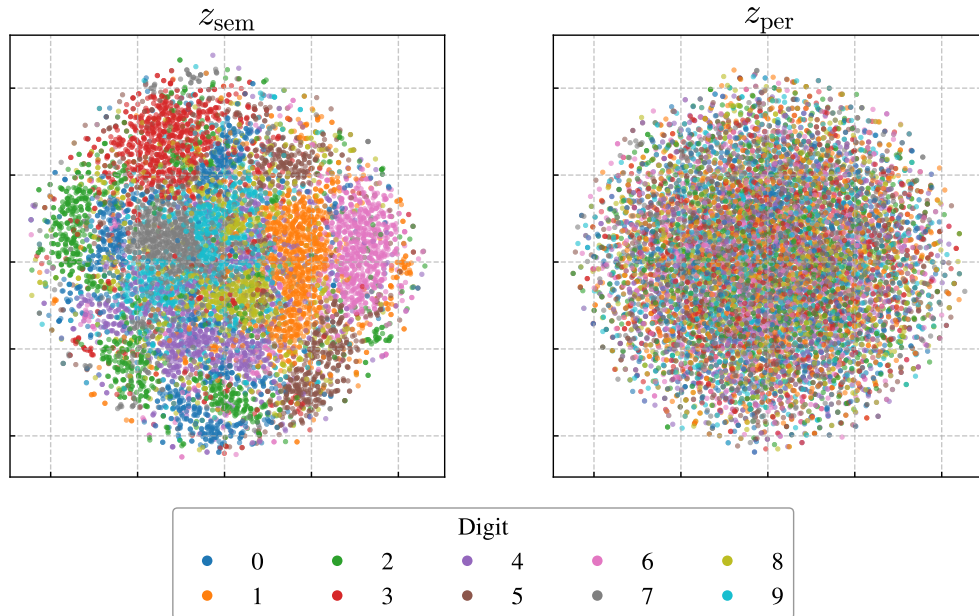


Figure 14: Neural entropy profiles for two image diffusion models trained on the MNIST and CIFAR-10 datasets. On the left is the entropy production rate, which is the time derivative of the neural entropy S_{NN}^{\bullet} , defined in Eq. (12). The sharp rise in entropy rate at early s is attributed to the low dimensionality of the data manifold. Notice how there is little difference between $S_{\text{NN}}^{X|Y}$ and S_{NN}^X , because the latter is *much* larger than $I(\mathbf{X}; \mathbf{Y})$ (cf. Eq. (11)). For MNIST, $S_{\text{NN}}^X = 2017.9$ and $I(\mathbf{X}; \mathbf{Y}) = 4.1$ (from Eq. (13)). For CIFAR-10, $S_{\text{NN}}^X = 7776.2$ and $I(\mathbf{X}; \mathbf{Y}) = 7.5404$. All measurements are in nats. The dashed red line indicates the partitioning of the denoising loss into semantic and perceptual pieces for the experiments in App. F.3.

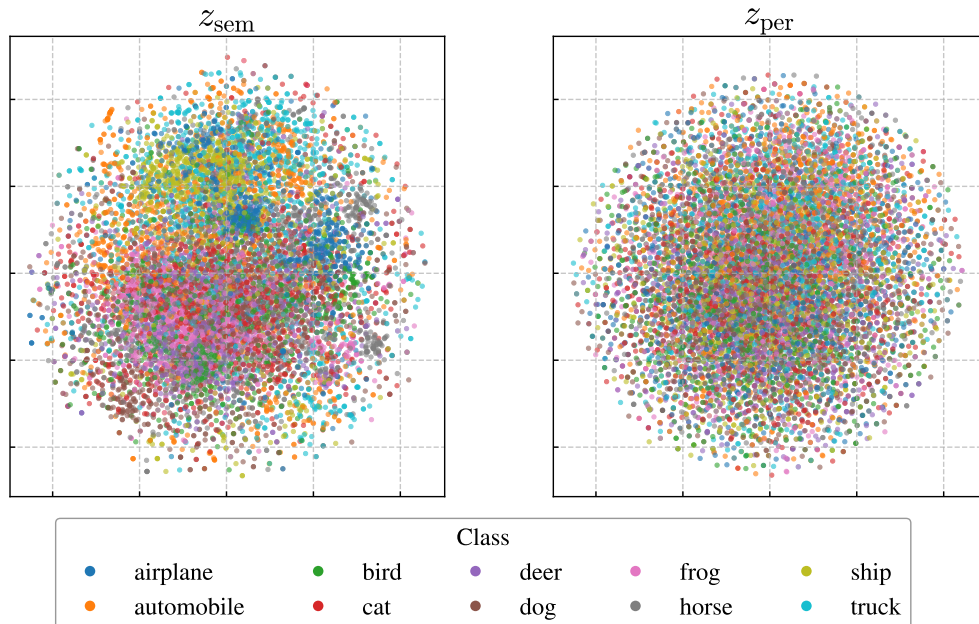
1674
 1675
 1676
 1677
 1678
 1679
 1680
 1681
 1682
 1683
 1684
 1685
 1686
 1687
 1688
 1689
 1690
 1691
 1692
 1693
 1694
 1695
 1696
 1697
 1698
 1699
 1700
 1701
 1702
 1703
 1704
 1705
 1706
 1707
 1708
 1709
 1710
 1711
 1712
 1713
 1714
 1715
 1716
 1717
 1718
 1719
 1720
 1721
 1722
 1723
 1724
 1725
 1726
 1727

2D t-SNE plot of MNIST latents



(a) A 2D t-SNE plot of the 20-dimensional latents z_{sem} and z_{per} produced by a DAE trained on MNIST digits. Information erased by the forward process up to $\tau = 0.1T$ is encoded in perceptual latent z_{per} , whereas all information beyond this point is captured by the semantic latent z_{sem} . Clusters of z_{sem} correspond to different MNIST digits. On the other hand, z_{per} shows little structure because the textural details of the images are very evenly distributed amongst all the digit classes.

2D t-SNE plot of CIFAR-10 latents



(b) t-SNE for CIFAR-10 latents. The more nuanced structure of z_{sem} reflects the far higher semantic variation between images in CIFAR-10. Both z_{sem} and z_{per} had $D_Z = 60$ dimensions.

Figure 15




Figure 16: Mutual information between the image labels Y and the corresponding semantic and perceptual latents, as a function of the partitioning time τ (cf. Eq. (74)). At early τ , the semantic latent is strongly correlated with the labels, whereas the perceptual details are completely independent. As τ increases, Z_{sem} becomes progressively irrelevant, whereas Z_{per} does the opposite—knowing enough small-to-medium details helps the model understand what the image is.

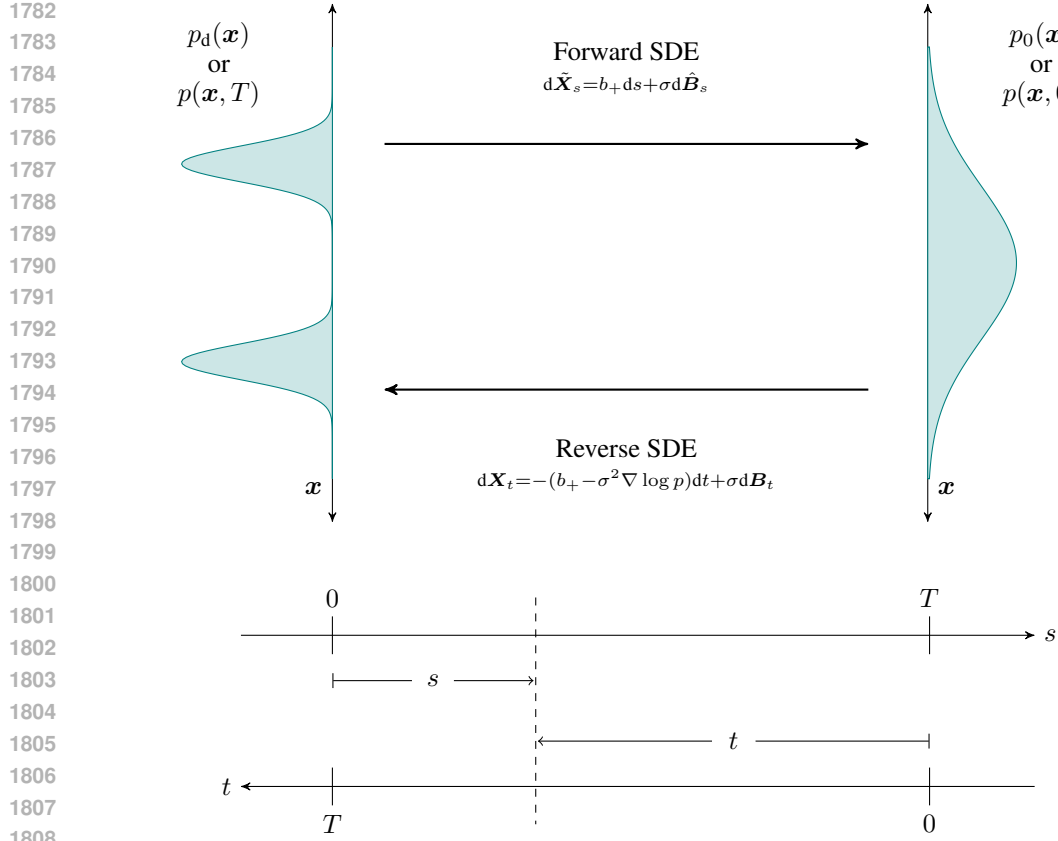
We begin by visualizing Z_{sem} and Z_{per} for DAE’s trained on MNIST and CIFAR-10 (see Fig. 15). We use $D_Z = 20$ for the former and $D_Z = 60$ for the latter, for both semantic and perceptual latents. These are generated by separate convolutional encoders. Optionally, we can adjust the receptive field of Z_{sem} to be larger than that of Z_{per} by increasing the number of encoder layers, as we do. These are mapped to two-dimensional space using t-SNE (van der Maaten & Hinton, 2008). With $\tau = 0.1T$, we find that there is little to no structure in Z_{per} in either case, in agreement with our claim that images from different classes have similar small-scale details. On the other hand, clusters of Z_{sem} appear in the t-SNE plot, showing that class labels correspond to large-scale features robust to small perturbations.⁴

We can do better than inspect Z_{sem} and Z_{per} by eye. Recall that Eq. (13) can be used to estimate the mutual information between random variables. By training a small diffusion model on pairs $\{(z_{\bullet}, \mathbf{y})\}_{i=1}^N$, we can find the determine $I(Z_{\bullet}, Y)$ approximately. The results are plotted in Fig. 16 for a range of τ values. As expected, Z_{sem} is correlates well with Y at small τ , whereas Z_{per} is nearly independent of it. However, as τ is increased Z_{per} rapidly encodes class information. We speculate that the $X \rightarrow Z_{\text{per}}$ encoder can detect semantic meaning if it’s given sufficient information about the medium-scale features, since an image is the sum of its parts. Furthermore, if τ is not too close to $s = 0$, the encoder focuses effort on information that differentiates the images, and downplays the shared textural detail between them. This is why Fig. 13 showed *some* clustering in the DAE case—even with the large amount of small-scale information the diffusion decoder captures, the encoder is incentivized to construct latents that uniquely identify the images (cf. App. F.1). We also mention in passing that the cross-over phenomenon in Fig. 13 is reminiscent of the *critical windows* of feature emergence (Li & Chen, 2024).

G NOTATION

The natural logarithm is denoted by \log . In Sec. 2 we use the symbol H for Shannon entropy of a discrete random variable, and $I(X, Y)$ is in bits. Everywhere else we use the differential entropy $S := -\int p \log p$, and the mutual information is in nats. Scalars are written in plain letters, while boldface symbols such as $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ denote higher-dimensional random variables. We write \mathbf{x} for a realization of \mathbf{X} , with unsubscripted symbols always referring to the data distribution p_d . We also write $p_d(\mathbf{x}, \mathbf{y})$ for the joint data distribution.

⁴This is why the diffusion classifiers from Clark & Jaini (2023); Li et al. (2023) employ a denoising objective that significantly downweights the contributions from the earlier time steps. The popular practice of ‘variance-dropping’ also achieves a similar effect (see App. B.3).



1809 Figure 17: A schematic of the forward and reverse diffusion processes.

1810
1811
1812 We use the time variable s for the forward diffusion process, which runs from left ($s = 0$) to
1813 right ($s = T$) in Fig. 17. \hat{B}_s and B_t denote the Brownian motions associated with the forward
1814 and reverse/controlled SDEs, respectively. ∇ is the gradient with respect the spatial coordinates,
1815 and ∂_t, ∂_s are partial time derivatives. S_{tot} is the total entropy produced during forward diffusion,
1816 and is closely approximated by the neural entropy S_{NN} . The time-dependence of the entropies is
1817 implicit in most of the main text; S_{tot} and S_{NN} without the time argument should be understood as
1818 $S_{\text{tot}}(s = T) \equiv S_{\text{tot}}(T)$.

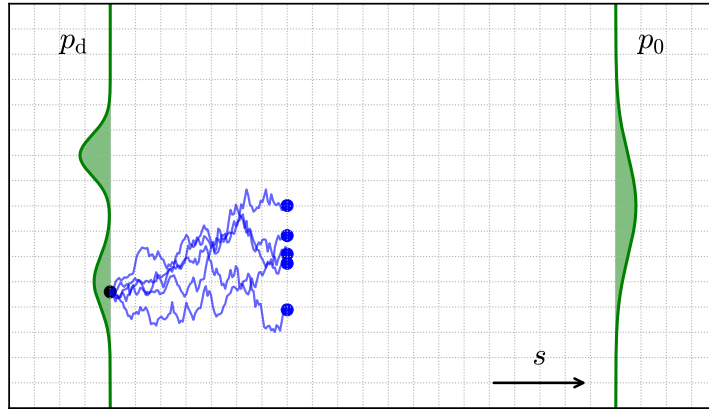
1819 The density $p(\tilde{x}_s, s)$ is the same as $p(x_t, t)$. That is, the symbol p is overloaded so we do not have to
1820 write $p(\cdot, s) = p(\cdot, T - t)$ everywhere. Throughout the paper, we set Boltzmann’s constant to unity,
1821 $k_B = 1$. p_d and p_0 denote the initial ($s = 0$) and final ($s = T$) densities for the forward process, and
1822 p_{eq} is its equilibrium state. Diffusion takes an infinite time to equilibrate, but we always take T to
1823 be large compared to the intrinsic time scale of the diffusion process.

1824 1825 H ILLUSTRATIONS

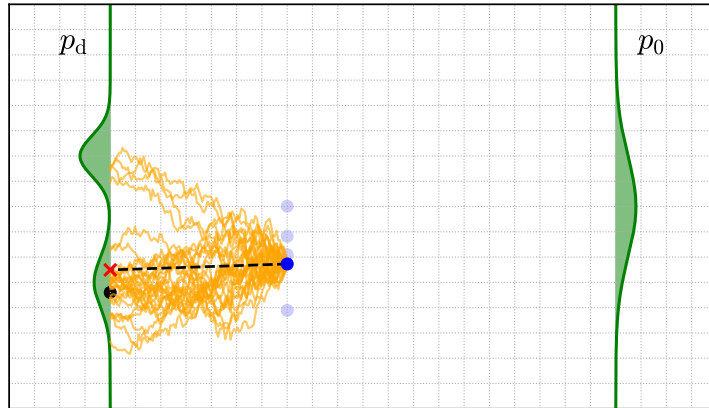
1826
1827 This section contains illustrations related to the discussion in App. C.

1828
1829
1830
1831
1832
1833
1834
1835


1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858
 1859
 1860
 1861
 1862
 1863
 1864
 1865
 1866
 1867
 1868
 1869
 1870
 1871
 1872
 1873
 1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889



(a) Forward diffusion from a test point x produces noisy samples \tilde{x}_s .




(b) $\hat{x}(\tilde{x}_s)$ is the average of all landings at $s = 0$ of reverse trajectories that start at \tilde{x}_s .




(c) Each \tilde{x}_s maps back to its own denoised mean $\hat{x}(\tilde{x}_s)$.

Figure 18: A breakdown of how $-\log p_d(x)$ is computed in in Eq. (45). The denoised means $\hat{x}(\tilde{x}_s)$ from many \tilde{x}_s give a sense of the regions in p_d that are most like x .


1890
 1891
 1892
 1893
 1894
 1895
 1896
 1897
 1898
 1899
 1900
 1901
 1902
 1903
 1904
 1905
 1906
 1907
 1908
 1909
 1910
 1911
 1912
 1913
 1914
 1915
 1916
 1917
 1918
 1919
 1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937
 1938
 1939
 1940
 1941
 1942
 1943



(a)



(b)



(c)

Figure 19: Iterative resolution of a sample from p_d . The blue points are forward diffused version of the black one. The red curve shows the density of denoised means computed from the blue points using the Miyasawa relation, Eq. (40). As we reverse diffuse in the t -direction, the red curve becomes more localized, gradually refining the generated sample toward its limiting value.