Deep Generative Modelling: A Comparative Review of VAEs, GANs, Normalizing Flows, Energy-Based and Autoregressive Models

Sam Bond-Taylor, Adam Leach, Yang Long, and Chris G. Willcocks

Abstract—Deep generative models are a class of techniques that train deep neural networks to model the distribution of training samples. Research has fragmented into various interconnected approaches, each of which make trade-offs including run-time, diversity, and architectural restrictions. In particular, this compendium covers energy-based models, variational autoencoders, generative adversarial networks, autoregressive models, normalizing flows, in addition to numerous hybrid approaches. These techniques are compared and contrasted, explaining the premises behind each and how they are interrelated, while reviewing current state-of-the-art advances and implementations.

Index Terms—Deep learning, generative models, energy-based models, variational autoencoders, generative adversarial networks, autoregressive models, normalizing flows

1 Introduction

Generative modelling using neural networks has its origins in the 1980s with aims to learn about data with no supervision, potentially providing benefits for standard classification tasks; collecting training data for unsupervised learning is naturally much lower effort and cheaper than collecting labelled data but there is considerable information still available making it clear that generative models can be beneficial for a wide variety of applications.

Beyond this, generative modelling has numerous direct applications including image synthesis: super-resolution, text-to-image and image-to-image conversion, inpainting, attribute manipulation, pose estimation; video: synthesis and retargeting; audio: speech and music synthesis; text: summarisation and translation; reinforcement learning; computer graphics: rendering, texture generation, character movement, liquid simulation; medical: drug synthesis, modality conversion; and out-of-distribution detection.

The central idea of generative modelling stems around training a generative model whose samples $\tilde{\boldsymbol{x}} \sim p_\theta(\tilde{\boldsymbol{x}})$ come from the same distribution as the training data distribution, $\boldsymbol{x} \sim p_d(\boldsymbol{x})$. Early neural generative models, energy-based models achieved this by defining an energy function on data points proportional to likelihood, however, these struggled to scale to complex high dimensional data such as natural images, and require Markov Chain Monte Carlo (MCMC) sampling during both training and inference, a slow iterative process. In recent years there has been renewed interest in generative models driven by the advent of large freely available datasets as well as advances in both general deep learning architectures and generative models, breaking new ground in terms of visual fidelity and sampling speed. In many cases, this has been achieved using latent variables z which are easy to sample from and/or calculate the density of, instead learning $p(\boldsymbol{x}, z)$; this requires marginalisation over the unobserved latent variables, however in general, this is intractable. Generative models therefore typically make trade-offs in execution time, architecture, or optimise proxy functions. Choosing what to optimise for has implications for sample quality, with direct likelihood optimisation often leading to worse sample quality than alternatives.

Interrelated with generative models is the field of self-supervised learning where the focus is on learning good intermediate representations that can be used for downstream tasks without supervision [105]. As such, generative models can in general also be considered self-supervised, however, not all self-supervised models are generative models. Types of self-supervised objectives include auxiliary classification losses such as predicting the rotation of inputs, masked losses where the model must predict the true value of some inputs which have been masked out, and contrastive losses which learn an embedding space where similar data points are close and different points are far apart.

There exists a variety of survey papers focusing on particular generative models such as normalizing flows [125], [175], generative adversarial networks [70], [249], and energy-based models [202], however, naturally these dive into the intricacies of their respective method rather than comparing with other methods; additionally, some focus on applications rather than theory. While there exists a recent

• *The authors are with the Department of Computer Science, Durham University, Durham DH1 3LE, U.K. E-mail: {samuel.e.bond-taylor, adam.leach, yang.long, christopher.g.willcocks}@durham.ac.uk.*

Manuscript received 14 April 2021; revised 5 July 2021; accepted 14 September 2021. Date of publication 30 September 2021; date of current version 3 October 2022.
This work was supported by MRC Innovation Fellowship with ref MR/S003916/1.
(Corresponding author: Sam Bond-Taylor.)
Recommended for acceptance by O. Winther.
Digital Object Identifier no. 10.1109/TPAMI.2021.3116668

TABLE 1
Comparison Between Deep Generative Models in Terms of Training and Test Speed, Parameter Efficiency, Sample Quality, Sample Diversity, and Ability to Scale to High Resolution Data. Quantitative Evaluation is Reported on the CIFAR-10 Dataset [127] in Terms of Fréchet Inception Distance (FID) and Negative Log-Likelihood (NLL) in Bits-Per-Dimension (BPD).

Method	Train Speed	Sample Speed	Num. Params.	Resolution Scaling	Free-form Jacobian	Exact Density	FID	NLL (in BPD)
Generative Adversarial Networks								
DCGAN [181]	★★★★★	★★★★★	★★★★★	★★★☆☆	✓	✗	37.11	-
ProGAN [114]	★★★★☆	★★★★★	★★★★☆	★★★★★	✓	✗	15.52	-
BigGAN [19]	★★★★☆	★★★★★	★★★★★	★★★★☆	✓	✗	14.73	-
StyleGAN2 + ADA [115]	★☆☆☆☆	★★★★★	★★★☆☆	★★★★★	✓	✗	2.42	-
Energy Based Models								
IGEBM [46]	★★★☆☆	★★☆☆☆	★★★★★	★★★☆☆	✓	✗	37.9	-
Denoising Diffusion [87]	★★★★★	★★☆☆☆	★★★☆☆	★★★★☆	✓	(✓)	3.17	≤ 3.75
DDPM++ Continuous [205]	★★★★★	★★☆☆☆	★★★☆☆	★★★★☆	✓	(✓)	2.20	-
Flow Contrastive (EBM) [55]	★☆☆☆☆	★★☆☆☆	★★☆☆☆	★★☆☆☆	✓	✗	37.30	≈ 3.27
VAEBM [246]	★★☆☆☆	★★★★☆	★☆☆☆☆	★★★★☆	✓	✗	12.19	-
Variational Autoencoders								
Convolutional VAE [123]	★★★★★	★★★★★	★★★★★	★★☆☆☆	✓	(✓)	106.37	≤ 4.54
Variational Lossy AE [29]	★★★☆☆	★☆☆☆☆	★☆☆☆☆	★★☆☆☆	✗	(✓)	-	≤ 2.95
VQ-VAE [183], [234]	★★☆☆☆	★★★☆☆	★☆☆☆☆	★★★★★	✗	(✓)	-	≤ 4.67
VD-VAE [31]	★☆☆☆☆	★★★★★	★★★☆☆	★★★★★	✓	(✓)	-	≤ 2.87
Autoregressive Models								
PixelRNN [233]	★★☆☆☆	★☆☆☆☆	★☆☆☆☆	★★★☆☆	✗	✓	-	3.00
Gated PixelCNN [232]	★★★☆☆	★☆☆☆☆	★☆☆☆☆	★★★☆☆	✗	✓	65.93	3.03
PixelIQN [173]	★★★☆☆	★☆☆☆☆	★☆☆☆☆	★★★☆☆	✗	✓	49.46	-
Sparse Trans. + DistAug [32], [110]	★★★★☆	★☆☆☆☆	★★★☆☆	★★★☆☆	✗	✓	14.74	2.66
Normalizing Flows								
RealNVP [43]	★★☆☆☆	★★★★★	★★★★★	★★★☆☆	✗	✓	-	3.49
GLOW [124]	★☆☆☆☆	★★★★★	★★☆☆☆	★★★★☆	✗	✓	45.99	3.35
FFJORD [62]	★☆☆☆☆	★★★☆☆	★★★★★	★★☆☆☆	✓	(✓)	-	3.40
Residual Flow [26]	★★☆☆☆	★★★★☆	★★★★☆	★★★★☆	✓	(✓)	46.37	3.28

survey on generative models as a whole [173], it is less broad, diving deeply into a few specific implementations.

This survey provides a comprehensive overview of generative modelling trends, introducing new readers to the field, comparing and contrasting so as to explain the modelling decisions behind each respective technique. Additionally, advances old and new are discussed in order to bring the reader up to date with current research. A specific focus on image models is taken reflecting the predominance in literature, however, concepts are often relevant across modalities. In particular, this survey covers energy-based models, unnormalised density models, variational autoencoders, variational approximation of a latent-based model's posterior, generative adversarial networks, two models set in a mini-max game, autoregressive models, model data decomposed as a product of conditional probabilities, and normalizing flows, exact likelihood models using invertible transformations. This breakdown is defined to closely match the typical divisions within research, however, numerous hybrid approaches exist that blur these lines, these are discussed in the most relevant section or both where suitable.

TABLE 2
Rules for the Star Ratings in Table 1

	1 Star	2 Stars	3 Stars	4 Stars	5 Stars
Training	>5 days	$\leq$5 days	$\leq$2 days	$\leq$1 days	$\leq \frac{1}{2}$ day
Sampling	AR	MCMC	Middle	$\leq$20 steps	1 step
Params	>120M	$\leq$120M	$\leq$60M	$\leq$30M	$\leq$10M
Resolution	<32	32	64 or 128	256 or 512	$\geq$1024

For a brief insight into the differences between architectures, we provide Table 1 which contrasts a diverse array of techniques. For the column "Exact Density", ✓ represents tractable densities, (✓) approximate densities, and ✗ intractable densities. On a number properties assessed we use a star system to allow easy comparisons, with rules defined in Table 2 based on CIFAR-10. In particular, we acknowledge that ranking measures such as training speed in days can be considered anecdotal since it is dependent on the year and compute available. Nevertheless, this allows a comparison based on properties such as stability and convergence rates which cannot be easily judged, for instance, by simply looking at number of function evaluations per iteration.

2 Energy-Based Models

Energy-based models (EBMs) [132] are based on the observation that any probability density function $p(\boldsymbol{x})$ for $\boldsymbol{x} \in \mathbb{R}^D$ can be expressed in terms of an energy function $E(\boldsymbol{x}) : \mathbb{R}^D \to \mathbb{R}$ which associates realistic points with low values and unrealistic points with high values

$$p(\boldsymbol{x}) = \frac{e^{-E(\boldsymbol{x})}}{\int_{\tilde{\boldsymbol{x}} \in \mathcal{X}} e^{-E(\tilde{\boldsymbol{x}})}}. \tag{1}$$

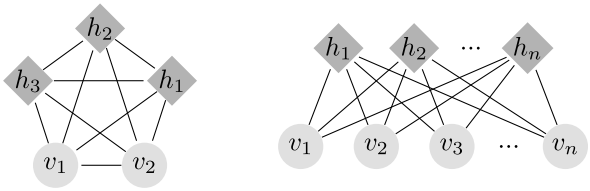

(a) Boltzmann machine. (b) Restricted Boltzmann machine.

Fig. 1. Restricted Boltzmann machines [82] have restricted architectures to allow faster sampling than Boltzmann machines [83].

Modelling data in such a way offers a number of perks, namely the simplicity and stability associating with training a single model; utilising a shared set of features thereby minimising required parameters; and the lack of any prior assumptions eliminates related bottlenecks [45]. Despite these benefits, scaling to high dimensional data is difficult, however, recent advances have made substantial strides.

A key issue with EBMs is how to optimise them; since the denominator in Eq. (1) is intractable for most models, a popular proxy objective is contrastive divergence where energy values of data samples are 'pushed' down, while samples from the energy distribution are 'pushed' up. Formally, the gradient of the negative log-likelihood loss $\mathcal{L}(\theta) = \mathbb{E}_{\boldsymbol{x}\sim p_d}[-\ln p_\theta(\boldsymbol{x})]$ has been shown to approximately demonstrate the following property [22], [206],

$$\nabla_\theta \mathcal{L} = \mathbb{E}_{\boldsymbol{x}^+\sim p_d}[\nabla_\theta E_\theta(\boldsymbol{x}^+)] - \mathbb{E}_{\boldsymbol{x}^-\sim p_\theta}[\nabla_\theta E_\theta(\boldsymbol{x}^-)], \tag{2}$$

where $\boldsymbol{x}^- \sim p_\theta$ is a sample from the EBM found through a Markov Chain Monte Carlo (MCMC) generating procedure.

2.1 Early Energy-Based Models

Before moving to recent advances, we start with some of the earliest neural generative models.

2.1.1 Boltzmann Machines

A Boltzmann machine [82] is a fully connected undirected network of binary neurons (Fig. 1) that are turned on with probability determined by a weighted sum of their inputs i.e. for some state s_i, $p(s_i = 1) = \sigma(\sum_j w_{i,j} s_j)$. The neurons can be divided into visible $\boldsymbol{v} \in \{0,1\}^D$ units, those which are set by inputs to the model, and hidden $\boldsymbol{h} \in \{0,1\}^P$ units, all other neurons. The energy of the state $\{\boldsymbol{v}, \boldsymbol{h}\}$ is defined (without biases for succinctness) as

$$E_\theta(\boldsymbol{v}, \boldsymbol{h}) = -\frac{1}{2}\boldsymbol{v}^T \boldsymbol{L}\boldsymbol{v} - \frac{1}{2}\boldsymbol{h}^T \boldsymbol{J}\boldsymbol{h} - \frac{1}{2}\boldsymbol{v}^T \boldsymbol{W}\boldsymbol{h}, \tag{3}$$

where $\boldsymbol{W}$, $\boldsymbol{L}$, and $\boldsymbol{J}$ are symmetrical learned weight matrices. In order to train Boltzmann machines via contrastive divergence, equilibrium states are found via Gibbs sampling, however, this takes an exponential amount of time in the number of hidden units making scaling impractical.

2.1.2 Restricted Boltzmann Machines

Many of the issues associated with Boltzmann machines can be overcome by restricting their connectivity. One approach, known as the restricted Boltzmann machine (RBM) [83] is to remove connections between units in the same group (Fig. 1b), allowing exact calculation of hidden units. Although obtaining negative samples still requires Gibbs sampling, it can be parallelised and in practice a single step is sufficient if $\boldsymbol{v}$ is initially sampled from the dataset [83].

By stacking RBMs, using features from lower down as inputs for the next layer, more powerful functions can be learned; these models are known as deep belief networks [84]. Training an entire model at once is intractable so instead they are trained greedily layer by layer, composing densities thus improving the approximation of $p(\boldsymbol{v})$.

2.2 Deep EBMs via Contrastive Divergence

To train more powerful architectures through contrastive divergence, one must be able to efficiently sample from p_θ. Specifically, we would like to model high dimensional data using an energy function with a deep neural network, taking advantage of recent advances in discriminative models [251]. MCMC methods such as random walk and Gibbs sampling [84], when applied to high dimensional data, have long mixing times, making them impractical. A number of recent approaches [45], [247] have advocated the use of stochastic gradient Langevin dynamics [186], [243] which permits sampling through the following iterative process,

$$\boldsymbol{x}_0 \sim p_0(\boldsymbol{x}) \qquad \boldsymbol{x}_{i+1} = \boldsymbol{x}_i - \frac{\alpha}{2}\frac{\partial E_\theta(\boldsymbol{x}_i)}{\partial \boldsymbol{x}_i} + \boldsymbol{\epsilon}, \tag{4}$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \alpha \boldsymbol{I})$, $p_0(\boldsymbol{x})$ is typically a uniform distribution over the input domain and α is the step size. As the number of updates $N \to \infty$ and $\alpha \to 0$, the distribution of samples converges to p_θ [243]; however, α and $\boldsymbol{\epsilon}$ are often tweaked independently to speed up training.

While Langevin MCMC is more practical than other approaches, sampling still requires a large number of steps. One solution is to use persistent contrastive divergence [45], [214] where a replay buffer stores previously generated samples that are randomly reset to noise; this allows samples to be continually refined with a relatively small number of steps while maintaining diversity. Short-run MCMC [165] which samples using as few as 100 update steps from noise has also been used to train deep EBMs, however, since the number of steps is so small, samples are not truly from the correct probability density. Nevertheless, there are other advantages such as allowing image interpolation and reconstruction (since short-run MCMC does not mix) [166]. Other approaches include initialising MCMC chains with data points [247] and samples from an implicit generative model [246], as well as adversarially training an implicit generative model, mitigating mode collapse somewhat by maximising its entropy [65], [120], [129]. Improved/augmented MCMC samplers with neural networks can also improve the efficiency of sampling [62], [88], [134], [199], [215].

One application of EBMs of this form comes by using standard classifier architectures, $f_\theta : \mathbb{R}^D \to \mathbb{R}^K$, which map data points to logits used by a softmax function to compute $p_\theta(y|\boldsymbol{x})$. By marginalising out y, these logits can be used to define an energy model that can be simultaneously trained as both a generative and classification model [64],

$$p_\theta(\boldsymbol{x}) = \sum_y p_\theta(\boldsymbol{x}, y) = \frac{\sum_y \exp(f_\theta(\boldsymbol{x}[y]))}{Z(\theta)}, \tag{5a}$$

$$E_\theta(\boldsymbol{x}) = -\ln \sum_y \exp(f_\theta(\boldsymbol{x}[y])). \tag{5b}$$

2.3 Score Matching and Denoising Diffusion

Although Langevin MCMC has allowed EBMs to scale to high dimensional data, training times are still slow due to the need to sample from the model distribution, additionally, the finite nature of the sampling process means that samples can be arbitrarily far away from the model's distribution [63]. An alternative approach is score matching [100] which is based on the idea of minimising the difference between the derivatives of the data and model's log-density functions; the score function is defined as $s(\boldsymbol{x}) = \nabla_{\boldsymbol{x}} \ln p(\boldsymbol{x})$ which does not depend on the intractable denominator and can therefore be applied to build an energy model [207] by minimising the Fisher divergence between p_d and p_θ,

$$\mathcal{L} = \frac{1}{2}\mathbb{E}_{p_d(\boldsymbol{x})}[\| \, s_\theta(\boldsymbol{x}) - s_d(\boldsymbol{x})\|_2^2], \tag{6}$$

however, the score function of data is usually not available. Various methods exist to estimate the score function including spectral approximation [194], sliced score matching [201], finite difference score matching [174], and notably denoising score matching [237] which allows the score to be approximated using corrupted data samples $q(\tilde{\boldsymbol{x}}|\boldsymbol{x})$. In particular, when $q = \mathcal{N}(\tilde{\boldsymbol{x}}|\boldsymbol{x}, \sigma^2\boldsymbol{I})$, Eq. (6) simplifies to

$$\mathcal{L} = \frac{1}{2}\mathbb{E}_{p_d(\boldsymbol{x})}\mathbb{E}_{\tilde{\boldsymbol{x}}\sim\mathcal{N}(\boldsymbol{x},\sigma^2\boldsymbol{I})}\left[\left\| s_\theta(\tilde{\boldsymbol{x}}) + \frac{\tilde{\boldsymbol{x}} - \boldsymbol{x}}{\sigma^2}\right\|_2^2\right]. \tag{7}$$

That is, s_θ learns to estimate the noise thereby allowing it to be used as a generative model [190], [200]. Since the Langevin update step uses $\nabla_{\boldsymbol{x}} \ln p(\boldsymbol{x})$ it is possible to sample from a score matching model using Langevin dynamics [224]. This is only possible, however, when trained over a large variety of noise levels so that $\tilde{\boldsymbol{x}}$ covers the whole space.

2.3.1 Denoising Diffusion Probabilistic Models

Closely related are diffusion models [1], [11], [86], [197] which gradually destroy data $\boldsymbol{x}_0$ by adding noise over a fixed number of steps T using a noise schedule $\beta_{1:T}$ determined so that $\boldsymbol{x}_T$ is approximately normally distributed. The forward process is defined by a discrete Markov chain,

$$q(\boldsymbol{x}_{1:T}|\boldsymbol{x}_0) = \prod_{t=1}^{T} q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}), \tag{8a}$$

$$q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}) = \mathcal{N}(\boldsymbol{x}_t; \sqrt{1-\beta_t}\boldsymbol{x}_{t-1}, \beta_t\boldsymbol{I}). \tag{8b}$$

The parameterised reverse process is trained to gradually remove noise, i.e. approximate $p_\theta(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t)$, by optimising a re-weighted variant of the ELBO, similar to Eq. (7).

Diffusion models have also been applied to categorical data; multinomial diffusions [92] define a forward process where each discrete variable switches randomly to a different value and the reverse process is trained to approximate the noise. Self-supervised language models such as BERT [40] have similar training objectives: variables are randomly masked out and the model is trained to predict the original values; these models can be viewed as Markov random fields and sampled using Gibbs/Metropolis Hastings via iterative sampling of the masked distributions [60], [238].

2.3.2 Speeding up Sampling

Sampling from score-based models requires a large number of steps leading to various techniques being developed to reduce this. A simple approach is to skip steps at inference: cosine schedules [161] spend more time where larger visual changes are made reducing the impact of skipping; another approach is to use dynamic programming to find what steps should be taken to minimise ELBO based on a computation budget [241]. Taking the continuous time limit of a diffusion model results in a stochastic differential equation (SDE), numerical solvers can therefore be used, reducing the number of steps required [107], [204]. Another proposed approach is to model noisy data points as $q(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t, \boldsymbol{x}_0)$, allowing the generative process to skip some steps using its approximation of end samples $\boldsymbol{x}_0$ [198].

2.4 Correcting Implicit Generative Models

While EBMs offer powerful representation ability due to unnormalized likelihoods, they can suffer from high variance training, long training and sampling times, and struggle to support the entire data space. In this section, a number of hybrid approaches are discussed which address these issues.

2.4.1 Exponential Tilting

To eliminate the need for an EBM to support the entire space, an EBM can instead be used to correct samples from an implicit generative network, simplifying the function to learn and allowing easier sampling. This procedure, referred to as exponentially tilting an implicit model, is defined as

$$p_{\theta,\phi}(\boldsymbol{x}) = \frac{1}{Z_{\theta,\phi}} q_\phi(\boldsymbol{x}) e^{-E_\theta(\boldsymbol{x})}. \tag{9}$$

By parameterising $q_\phi(\boldsymbol{x})$ as a latent variable model such as a normalizing flow [3], [164] or VAE generator [245], MCMC sampling can be performed in the latent space rather than the data space. Since the latent space is much simpler, and often uni-modal, MCMC mixes much more effectively. This limits the freedom of the model, however, leading some to jointly sample in latent and data space [3], [245].

2.4.2 Noise Contrastive Estimation

Noise contrastive estimation [51], [74] transforms EBM training into a classification problem using a noise distribution $q_\phi(\boldsymbol{x})$ by optimising the loss function,

$$\mathbb{E}_{p_d}\left[\ln \frac{p_\theta(\boldsymbol{x})}{p_\theta(\boldsymbol{x}) + q_\phi(\boldsymbol{x})}\right] + \mathbb{E}_{q_\phi}\left[\ln \frac{q_\phi(\boldsymbol{x})}{p_\theta(\boldsymbol{x}) + q_\phi(\boldsymbol{x})}\right], \tag{10}$$

where $p_\theta(\boldsymbol{x}) = e^{E_\theta(\boldsymbol{x})-c}$. This approach can be used to train a correction via exponential tilting [164], but can also be used to directly train an EBM and normalizing flow [54]. Eq. (10) is equivalent to GAN Equation (18), however, training formulations differ, with noise contrastive estimation explicitly modelling likelihood ratios.

2.5 Alternative Training Objectives

As aforementioned, energy models trained with contrastive divergence approximately maximises the likelihood of the data; likelihood however does not correlate directly with sample quality [213]. Training EBMs with arbitrary f-divergences is possible, yielding improved FID scores [250].

Since score estimates have high variance, the Stein discrepancy has been proposed as an alternative objective, requiring no sampling and more closely correlating with likelihood [63]. A middle ground between denoising score matching and contrastive divergence is diffusion recovery likelihood [11] which can be optimised via a sequence of denoising EBMs conditioned on increasingly noisy samples of the data, the conditional distributions being much easier to MCMC sample from than typical EBMs [55].

3 Variational Autoencoders

One of the key problems associated with energy-based models is that sampling is not straightforward and mixing can require a significant amount of time. To circumvent this issue, it would be beneficial to explicitly sample from the data distribution with a single network pass.

To this end, suppose we have a latent based model $p_\theta(\boldsymbol{x}|\boldsymbol{z})$ with prior $p_\theta(\boldsymbol{z})$ and posterior $p_\theta(\boldsymbol{z}|\boldsymbol{x})$; unfortunately optimising this model through maximum likelihood is intractable due to the integral in $p_\theta(\boldsymbol{x}) = \int_{\boldsymbol{z}} p_\theta(\boldsymbol{x}|\boldsymbol{z})p_\theta(\boldsymbol{z})d\boldsymbol{z}$. Instead, variational inference allows this problem to be reframed as an optimisation problem by introducing an approximation of the true intractable posterior $q_\phi(\boldsymbol{z}|\boldsymbol{x}) = \arg\min_q D_{KL}(q_\phi(\boldsymbol{z}|\boldsymbol{x})||p_\theta(\boldsymbol{z}|\boldsymbol{x}))$ that allows a tractable bound on $p_\theta(\boldsymbol{x})$ to be formed. In particular, variational autoencoders amortize the inference process, that is, approximate $q_\phi(\boldsymbol{z}|\boldsymbol{x})$ using a feedforward inference network allowing scaling to large datasets [122], [185]. From the definition of KL divergence we get

$$\begin{aligned} D_{KL}(q_\phi(\boldsymbol{z}|\boldsymbol{x})||p_\theta(\boldsymbol{z}|\boldsymbol{x})) &= \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}\left[\ln \frac{q_\phi(\boldsymbol{z}|\boldsymbol{x})}{p_\theta(\boldsymbol{z}|\boldsymbol{x})}\right] \\ &= \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln q_\phi(\boldsymbol{z}|\boldsymbol{x})] - \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln p_\theta(\boldsymbol{z},\boldsymbol{x})] \\ &\quad + \ln p_\theta(\boldsymbol{x}), \end{aligned} \tag{11}$$

which can be rearranged to find an alternative definition for $p_\theta(\boldsymbol{x})$ that does not require the knowledge of $p_\theta(\boldsymbol{z}|\boldsymbol{x})$

$$\begin{aligned} \ln p_\theta(\boldsymbol{x}) &= D_{KL}(q_\phi(\boldsymbol{z}|\boldsymbol{x})||p_\theta(\boldsymbol{z}|\boldsymbol{x})) - \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln q_\phi(\boldsymbol{z}|\boldsymbol{x})] \\ &\quad + \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln p_\theta(\boldsymbol{z},\boldsymbol{x})] \\ &\geq -\mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln q_\phi(\boldsymbol{z}|\boldsymbol{x})] + \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln p_\theta(\boldsymbol{z},\boldsymbol{x})] \\ &= -\mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln q_\phi(\boldsymbol{z}|\boldsymbol{x})] + \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln p_\theta(\boldsymbol{z})] \\ &\quad + \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln p_\theta(\boldsymbol{x}|\boldsymbol{z})] \\ &= -D_{KL}(q_\phi(\boldsymbol{z}|\boldsymbol{x})||p_\theta(\boldsymbol{z})) + \mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln p_\theta(\boldsymbol{x}|\boldsymbol{z})] \\ &\equiv \mathcal{L}(\theta,\phi;\boldsymbol{x}), \end{aligned} \tag{12}$$

where $\mathcal{L}$ is known as the evidence lower bound (ELBO) [108]. To optimise this bound with respect to θ and ϕ, gradients must be backpropagated through the stochastic sampling process $\tilde{\boldsymbol{z}} \sim q_\phi(\boldsymbol{z}|\boldsymbol{x})$. This is permitted by reparameterizing $\tilde{\boldsymbol{z}}$ using a differentiable function $g_\phi(\boldsymbol{\epsilon},\boldsymbol{x})$ of a noise variable $\boldsymbol{\epsilon}$: $\tilde{\boldsymbol{z}} = g_\phi(\boldsymbol{\epsilon},\boldsymbol{x})$ with $\boldsymbol{\epsilon} \sim p(\boldsymbol{\epsilon})$ [122], [185].

Monte Carlo gradient estimators can be used to approximate the expectations, however, this yields very high variance making it impractical. Alternatively, if $D_{KL}(q_\phi(\boldsymbol{z}|\boldsymbol{x})||p_\theta(\boldsymbol{z}))$ can be integrated analytically then the variance is manageable. A prior with such a property needs to be simple enough to sample from but also sufficiently flexible to match the true posterior; a common choice is a normally distributed prior with diagonal covariance, $\boldsymbol{z} \sim q_\phi(\boldsymbol{z}|\boldsymbol{x}) = \mathcal{N}(\boldsymbol{z};\boldsymbol{\mu},\boldsymbol{\sigma}^2\boldsymbol{I})$ with $\tilde{\boldsymbol{z}} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}$ and $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0},\boldsymbol{I})$. In this case, the loss simplifies to

$$\begin{aligned} \tilde{\mathcal{L}}_{VAE}(\theta,\phi;\boldsymbol{x}) \simeq \frac{1}{2}\sum_{j=1}^{J}\left(1 + \ln((\sigma^{(j)})^2) - (\mu^{(j)})^2 - (\sigma^{(j)})^2\right) \\ + \frac{1}{L}\sum_{l=1}^{L} \ln p_\theta(\boldsymbol{x}|\tilde{\boldsymbol{z}}_l). \end{aligned} \tag{13}$$

Despite success on small scale datasets, when applied to more complex datasets such as natural images, samples tend to be unrealistic and blurry [44]. This blurriness has been attributed to the maximum likelihood objective itself and MSE reconstruction loss, however, there is evidence that limited approximation of the true posterior is the root cause [258]; with MSE causing highly non-Gaussian posteriors. As such, the Gaussian posterior implies an overly simple model which, when unable to perfectly fit, maps multiple data points to the same encoding leading to averaging.

There are a number of other issues associated with limited posterior approximation, namely under-estimation of the variance of the posterior, resulting in poor predictions, and biases in the MAP estimates of model parameters [222]. Additionally, amortized inference leads to an amortization gap, the difference in ELBO for the amortized posterior and optimal approximate posterior [36]. Increasing the capacity of the encoder and decoder can reduce this gap by improving the posterior approximation and better fitting the choice of approximation respectively. Other proposed improvements include combining with adversarial training [97], [131], [149], improving the ELBO [20], as well as using different regularisation such as Wasserstein distance [216].

Reweighting the ELBO by multiplying D_{KL} with an extra hyperparameter β allows the capacity of the latent representation to be altered. When $\beta > 1$ a more disentangled representation is learned where each latent unit is responsible for a single generative factor [81]. This approach has been generalised, allowing more precise states in the compression-representation trade-off to be targeted [2].

3.1 Beyond Simple Priors

One approach to improve variational bounds and increase sample quality is to improve the priors used for instance by careful selection to the task or by increasing its complexity [89]. Complex priors can be learned by warping simple distributions and inducing variational dependencies between the latent variables: variational Gaussian processes permit this by forming an infinite ensemble of mean-field distributions [218]; EBMs and score matching can be used to model flexible priors [265], [227]; normalizing flows (see Section 6)

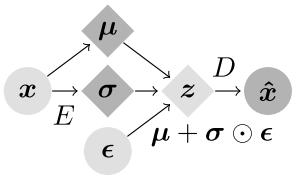

Fig. 2. Variational autoencoder [122], [185] with a normally distributed prior. ϵ is sampled from $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$.

transform distributions through a series of invertible parameterised functions [13], [61], [96], [124], [184], [189].

By rewriting the VAE training objective to have two regularisation terms [149],

$$\begin{aligned}\mathcal{L}(\theta, \phi; \boldsymbol{x}) = \mathbb{E}_{\boldsymbol{x}\sim q(\boldsymbol{x})}[\mathbb{E}_{q_\phi(\boldsymbol{z}|\boldsymbol{x})}[\ln p_\theta(\boldsymbol{x}|\boldsymbol{z})]] \\ + \mathbb{E}_{\boldsymbol{x}\sim q(\boldsymbol{x})}[\mathbb{H}[q_\phi(\boldsymbol{z}|\boldsymbol{x})]] - \mathbb{E}_{z\sim q(\boldsymbol{z})}[p_\theta(\boldsymbol{z})],\end{aligned} \tag{14}$$

the latter of which is the cross entropy between the aggregate posterior and the prior, the prior can be defined as the aggregate posterior, thus obtaining a rich multi-modal latent representation that combats inactive latent variables. Since the true aggregate posterior is intractable, VampPrior [217] approximates it for a set of pseudo-inputs, tensors with the same shape as data points learned during training. Exemplar VAEs [168] scale this approach up, using the full training set to approximate the aggregate posterior, by approximating the prior using k-nearest-neighbours. Alternatively, the aggregate posterior can be approximated with a learned prior; this has been achieved with a learned rejection sampling procedure that transforms a base distribution [7].

In some instances, it can be helpful to compress data to discrete latent representations [17], [110], however, gradients through discrete sampling procedures are ill-defined. The Gumbel-Softmax/Concrete distribution is a differentiable continuous approximation of a categorical distribution containing a temperature coefficient that converges to a discrete distribution in the limit [103], [147].

Alternatively, it has been argued that simple Gaussian priors are not a hindrance. When the data of dimension d lies on a sub-manifold of dimension r and $r < d$ then global VAE optimum exist that do no recover the data distribution, however, when $r = d$, global optimums do recover the data distribution; as such, 2 stage VAEs that first map data to latents of dimension r then use a second VAE to correct the learned density can better capture the data [37].

3.1.1 Hierarchical VAEs

Hierarchical VAEs build complex priors with multiple levels of latent variables, each conditionally dependent on the last, forming dependencies depthwise though the network,

$$p_\theta(\boldsymbol{z}) = p_\theta(\boldsymbol{z}_0)p_\theta(\boldsymbol{z}_1|\boldsymbol{z}_0)\cdots p_\theta(\boldsymbol{z}_N|\boldsymbol{z}_{<N}), \tag{15a}$$

$$q_\phi(\boldsymbol{z}|\boldsymbol{x}) = q_\phi(\boldsymbol{z}_0|\boldsymbol{x})q_\phi(\boldsymbol{z}_1|\boldsymbol{z}_0, \boldsymbol{x})\cdots q_\phi(\boldsymbol{z}_N|\boldsymbol{z}_{<N}, \boldsymbol{x}). \tag{15b}$$

Ladder VAEs [209] achieve this conditioning structure using a bidirectional inference network where a deterministic "bottom-up" pass generates features at various resolutions, then the latent variables are processed from top to bottom with the features shared (Fig. 3). Specifically, they model latents as normal distributions conditioned on the last latent,

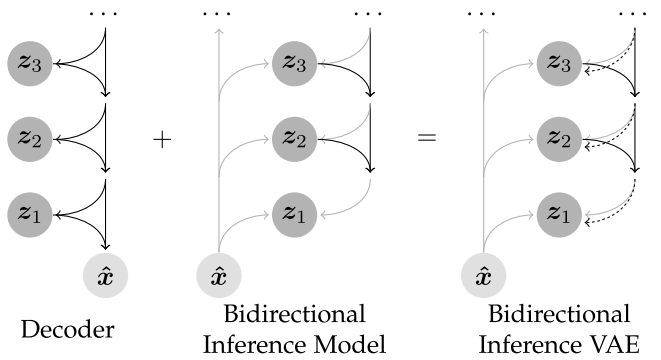

Fig. 3. A hierarchical VAE with bidirectional inference [125].

$$p_\theta(\boldsymbol{z}_i|\boldsymbol{z}_{i-1}) = \mathcal{N}(\boldsymbol{z}_i|\mu_{p,i}(\boldsymbol{z}_{i-1}), \sigma^2_{p,i}(\boldsymbol{z}_{i-1})). \tag{16}$$

By introducing skip connections around the stochastic sampling process, latents can be conditioned on all previously sampled latents [124], [146], [226]. Such an architecture generalises autoregressive models; inferring latents in parallel allows for significantly fewer steps compared to typical autoregressive models since many latents are statistically independent and allows different latent levels to correspond to global/local details depending on their depth. It has been argued that a single level of latents is sufficient since Gibbs sampling performed on that level can recover the data distribution [257]. Despite that, Gibbs sampling converges slowly, making hierarchical representations more efficient; in support of this, deeper hierarchical VAEs have been shown to improve likelihood, independent of capacity [30].

3.2 Regularised Autoencoders

Related to VAEs are regularised autoencoders (RAEs) which apply regularisation to the latent space of a deterministic autoencoder then subsequently train a density estimator on this space to obtain a complex prior [57]. Since the approximate posterior is a degenerate distribution, RAEs have little connection with variational inference. Vector Quantized-Variational Autoencoders (VQ-VAE) [181], [233] achieve this by training an autoencoder with a discrete latent space, then approximating encodings with an autoregressive model (see Section 5). The encoder's outputs are compared to a codebook of latent vectors and set to the code they are closest to; the gradient of this discretisation process is approximated using the straight through estimator [10]. Meanwhile, latent vectors in the codebook are moved closer to the encoder's outputs. To model larger images, hierarchy of codes have been applied [182], as well as adversarial learning to increase compression rate [52].

3.3 Data Modelling Distributions

Unlike energy-based models, VAEs must model an explicit density $p(\boldsymbol{x}|\boldsymbol{z})$. For efficient sampling, typically this distribution is decomposed as a product of independent simple distributions, allowing unrestricted architectures to be used to parameterise the chosen distributions. Common instances include modelling variables as Bernoulli [141], Gaussian [122], multinomial distributions, or as mixtures [188].

3.3.1 Autoregressive Decoders

To introduce dependencies between the output variables, numerous works have used powerful autoregressive networks [72]. While these approaches allow complex distributions to be learned, they increase the runtime and often suffer from posterior collapse since early in training the approximate posterior contains little knowledge about $\boldsymbol{x}$ meaning that it is easy to minimise D_{KL} which in turn reduces the gradient between the encoder and decoder making it difficult to escape this minima [17]; in fact, for a sufficiently powerful generative distribution, this can occur even at optimum solutions [28]. Various methods to prevent posterior collapse have been proposed: by restricting the autoregressive network's receptive field to a small window, it is forced to use latents to capture global structure [28]; a mutual information term can be added to the loss to encourage high correlation between $\boldsymbol{x}$ and $\boldsymbol{z}$ [259]; encouraging the posterior to be diverse by controlling its geometry to evenly covering the data space, redundancy is reduced and latents are encouraged to learn global structure [145].

3.4 Bridging Amortized and Stochastic Inference

While variational approaches offer substantial speedup over MCMC sampling, there is an inherent discrepancy between the true posterior and approximate posterior despite improvements in this field. To this end, a number of approaches have been proposed to find a middle ground, yielding improvements over amortized methods with lower costs than MCMC. Semi-amortised VAEs [121] use an encoder network followed by stochastic gradient descent on latents to improve the ELBO, however, this still relies on an inference network. The inference network can be removed by assigning latent vectors to data points, then optimising them with Langevin dynamics or gradient descent, during training; although this allows fast training, convergence for unseen samples is not guaranteed and there is still a large discrepancy between the true posterior and latent approximations due to lag in optimisation [15], [77]. Short-run MCMC has also been applied however it has poor mixing properties [167]. Gradient Origin Networks [16] replace the encoder with an empirical Bayes approximation of the posterior that only requires a single gradient step.

VAEBMs offer a different perspective, rather than performing latent MCMC sampling based on the ELBO, they use an auxiliary energy-based model to correct blurry VAE samples, with MCMC sampling performed in both the data space and latent space. This setup is defined by $h_{\phi,\theta}(\boldsymbol{x},\boldsymbol{z}) = \frac{1}{Z_{\phi,\theta}} p_\theta(\boldsymbol{z}) p_\theta(\boldsymbol{x}|\boldsymbol{z}) e^{-E_\phi(\boldsymbol{x})}$, where $p_\theta(\boldsymbol{z})p_\theta(\boldsymbol{x}|\boldsymbol{z})$ is the VAE, and $E_\phi(\boldsymbol{x})$ is the energy model. This, however, requires 2 stages of training to avoid calculating the gradient of the normalising constant $Z_{\phi,\theta}$, training only the VAE and fixing the VAE and training the EBM respectively.

4 Generative Adversarial Networks

Another approach at eliminating the Markov chains used in energy models is the generative adversarial network (GAN) [58]. GANs consist of two networks, a discriminator $D: \mathbb{R}^n \to [0,1]$ which estimates the probability that a sample comes from the data distribution $\boldsymbol{x} \sim p_d(\boldsymbol{x})$, and a generator

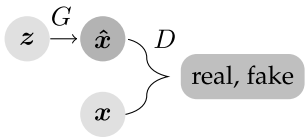

Fig. 4. Generative adversarial networks [58] set two networks in a game: D detects real from fake samples while G tricks D.

$G: \mathbb{R}^m \to \mathbb{R}^n$ which given a latent variable $\boldsymbol{z} \sim p_z(z)$, captures p_d by tricking the discriminator into thinking its samples are real. This is achieved through adversarial training of the networks: D is trained to correctly label training samples as real and samples from G as fake, while G is trained to minimise the probability that D classifies its samples as fake. This can be interpreted as D and G playing a minimax game, as with prior work [192], [193], optimising the value function $V(G,D)$,

$$\begin{aligned}\min_G \max_D V(G,D) = \mathbb{E}_{\boldsymbol{x}\sim p_d(\boldsymbol{x})}[\ln D(\boldsymbol{x})] \\ + \mathbb{E}_{\boldsymbol{z}\sim p_z(\boldsymbol{z})}[\ln(1 - D(G(\boldsymbol{z})))].\end{aligned} \tag{17}$$

For a fixed G, the objective for D can be reformulated as

$$\begin{aligned}\max_D V(G,D) &= \mathbb{E}_{\boldsymbol{x}\sim p_d}[\ln D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{x}\sim p_g}[\ln(1 - D(\boldsymbol{x}))] \\ &= \mathbb{E}_{\boldsymbol{x}\sim p_d}\left[\ln \frac{p_d(\boldsymbol{x})}{p_d(\boldsymbol{x}) + p_g(\boldsymbol{x})}\right] \\ &\quad + \mathbb{E}_{\boldsymbol{x}\sim p_g}\left[\ln \frac{p_g(\boldsymbol{x})}{p_d(\boldsymbol{x}) + p_g(\boldsymbol{x})}\right] \\ &= D_{KL}(p_d||\tfrac{1}{2}(p_d + p_g)) + D_{KL}(p_g||\tfrac{1}{2}(p_d + p_g)) + C.\end{aligned} \tag{18}$$

Therefore the loss is equivalent to the Jensen-Shannon divergence between the generative distribution p_g and the data distribution p_d and thus with sufficient capacity, the generator can recover the data distribution. The use of symmetric JS-divergence is well behaved when both distributions are small unlike the asymmetric KL-divergence used in maximum likelihood models. Additionally, it has been suggested that reverse KL-divergence, $D_{KL}(p_g||p_d)$, is a better measure for training generative models than normal KL-divergence, $D_{KL}(p_d||p_g)$, since it minimises $\mathbb{E}_{\boldsymbol{x}\sim p_g}[\ln p_d(\boldsymbol{x})]$ [99]; while reverse KL-divergence is not a viable objective function, JS-divergence is and behaves more like reverse KL-divergence than KL-divergence alone. With that said, JS-divergence is not perfect; if 0 mass is associated with a data sample in a maximum likelihood model, KL-divergence is driven to infinity, whereas this can happen with no consequence in a GAN.

4.1 Stabilising Training

The adversarial nature of GANs makes them notoriously difficult to train [4]; Nash equilibrium is hard to achieve [187] since non-cooperation cannot guarantee convergence, thus training often results in oscillations of increasing amplitude. As the discriminator improves, gradients passed to the generator vanish, accelerating this problem; on the other hand, if the discriminator remains poor, the generator does not receive useful gradients. Another problem is mode

Name	Discriminator Loss	Generator Loss
NSGAN [59]	$-\mathbb{E}[\ln(\sigma(D(\boldsymbol{x})))] - \mathbb{E}[\ln(1-\sigma(D(G(\boldsymbol{z}))))]$	$-\mathbb{E}[\ln(\sigma(D(G(\boldsymbol{z}))))]$
WGAN [5]	$\mathbb{E}[D(\boldsymbol{x})] - \mathbb{E}[D(G(\boldsymbol{z}))]$	$\mathbb{E}[D(G(\boldsymbol{z}))]$
LSGAN [151]	$\mathbb{E}[(D(\boldsymbol{x})-1)^2] + \mathbb{E}[D(G(\boldsymbol{z}))^2]$	$\mathbb{E}[(D(G(\boldsymbol{z}))-1)^2]$
Hinge [136]	$\mathbb{E}[\min(0, D(\boldsymbol{x})-1)] - \mathbb{E}[\max(0, 1+D(G(\boldsymbol{z})))]$	$-\mathbb{E}[D(G(\boldsymbol{z}))]$
EBGAN [257]	$D(\boldsymbol{x}) + \max(0, m - D(G(\boldsymbol{z})))$	$D(G(\boldsymbol{z}))$
RSGAN [107]	$\mathbb{E}[\ln(\sigma(D(\boldsymbol{x}) - D(G(\boldsymbol{z}))))]$	$\mathbb{E}[\ln(\sigma(-D(G(\boldsymbol{z})) - D(\boldsymbol{x})))]$

(a) GAN losses.

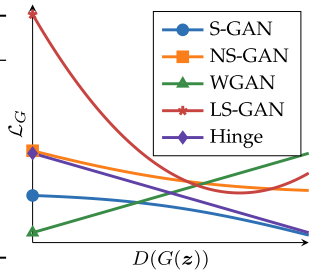

(b) Generator loss functions.

Fig. 5. A comparison of popular losses used to train GANs. (a) Respective losses for discriminator/generator. (b) Plots of generator losses with respect to discriminator output. Notably, NS-GAN's gradient disappears as discriminator gets better.

collapse, where one network gets stuck in a bad local minima and only a small subset of the data distribution is learned. The discriminator can also jump between modes resulting in catastrophic forgetting, where previously learned knowledge is forgotten when learning something new [211]. This section explores proposed solutions to these problems.

4.1.1 Loss Functions

Since the cause of many of these issues can be linked with the use of JS-divergence, other loss functions have been proposed that minimise other statistical distances; in general, any f-divergence can be used to train GANs [169]. One notable example is the Wasserstein distance which intuitively indicates how much "mass" must be moved to transform one distribution into another. Wasserstein distance is defined formally in Eq. (19a), which by the Kantorovich-Rubinstein duality is equivalent to Eq. (19b) [236]

$$W(p_d, p_g) = \inf_{\gamma \in \prod(p_d, p_g)} \mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})\sim\gamma}[\|\boldsymbol{x}-\boldsymbol{y}\|], \tag{19a}$$

$$W(p_d, p_g) = \sup_{\|D\|_L \leq 1} \mathbb{E}_{\boldsymbol{x}\sim p_d}[D(\boldsymbol{x})] - \mathbb{E}_{\boldsymbol{x}\sim p_g}[D(\boldsymbol{x})], \tag{19b}$$

where the supremum is taken over all 1-Lipschitz functions, that is, f such that for all x_1 and x_2, $\|f(x_1) - f(x_2)\|_2 \leq \|x_1 - x_2\|_2$. Optimising Wasserstein distance, as described in Table 5a, offers linear gradients thus eliminating the vanishing gradients problem (see Fig. 5). Moreover, Wasserstein distance is also equivalent to minimising reverse KL-divergence [156], offers improved stability, and allows training to optimality. Numerous approaches to enforce 1-Lipschitz continuity have been proposed: weight clipping [5] invalidates gradients making optimisation difficult; applying a gradient penalty within the loss is heavily dependent on the support of the generative distribution and computation with finite samples makes application to the entire space intractable [71]; spectral normalisation (discussed below) applies global regularisation by estimating the singular values of parameters. Other popular loss functions include least squares GAN, hinge loss, energy-based GAN, and relativistic GAN (detailed in Table 5a).

The catastrophic forgetting problem can be mitigated by conditioning the GAN on class information, encouraging more stable representations [18], [155], [253]. Nevertheless, labelled data, if available, only covers limited abstractions. Self-supervision achieves the same goal by training the discriminator on an auxiliary classification task based solely on the unsupervised data. Proposed approaches are based on randomly rotating inputs to the discriminator, which learns to identify the angle rotated separately to the standard real/fake classification [27]. Extensions include training the discriminator to jointly determine rotation and real/fake to provide better feedback [221], and training the generator to trick the discriminator at both the real/fake and classification tasks [221]. A more explicit approach is to model the generator with a normalizing flow, avoiding collapse by jointly optimising the GAN and likelihood objectives [69].

4.1.2 Spectral Normalisation

Spectral normalisation [156] is a technique to make a function globally 1-Lipschitz utilising the observation that the Lipschitz constant of a linear function is its largest singular value (spectral norm). The spectral norm of a matrix $\boldsymbol{A}$ is

$$SN(\boldsymbol{A}) := \max_{\boldsymbol{h}:\boldsymbol{h}\neq\boldsymbol{0}} \frac{\|\boldsymbol{A}\boldsymbol{h}\|_2}{\|\boldsymbol{h}\|_2} = \max_{\|\boldsymbol{h}\|_2 \leq 1} \|\boldsymbol{A}\boldsymbol{h}\|_2, \tag{20}$$

thus a weight matrix $\boldsymbol{W}$ is normalised to be 1-Lipschitz by replacing the weights with $\boldsymbol{W}_{SN} := \frac{\boldsymbol{W}}{SN(\boldsymbol{W})}$. Rather than using singular value decomposition to compute the norm, the power iteration method is used; for randomly initialised vectors $\boldsymbol{v} \in \mathbb{R}^n$ and $\boldsymbol{u} \in \mathbb{R}^m$, the procedure is

$$\boldsymbol{u}_{t+1} = \boldsymbol{W}\boldsymbol{v}_t, \ \boldsymbol{v}_{t+1} = \boldsymbol{W}^T\boldsymbol{u}_{t+1}, \ SN(\boldsymbol{W}) \approx \boldsymbol{u}^T\boldsymbol{W}\boldsymbol{v}. \tag{21}$$

Since weights change only marginally with each optimisation step, a single power iteration step per global optimisation step is sufficient to keep $\boldsymbol{v}$ and $\boldsymbol{u}$ close to their targets.

As aforementioned, enforcing the discriminator to be 1-Lipschitz is essential for WGANs, however, spectral normalisation has been found to dramatically improve sample quality and allow scaling to datasets with thousands of classes across a variety of loss functions [18], [156]. Spectral collapse, has been linked to discriminator overfitting when spectral norms of layers explode [18] as well as mode collapse when spectral norms fall in value significantly [138]. Additionally, regularising the discriminator in this manner helps balance the two networks, reducing the number of discriminator update steps required [18], [253].

4.1.3 Data Augmentation

Augmenting training data to increase the quantity of training data is often common practice; when training GANs the types of augmentations permitted are limited to more simple augmentations such as cropping and flipping to prevent the generator from creating undesired artefacts. Several approaches independently proposed applying augmentations to all discriminator inputs, allowing more substantial augmentations to be used [114], [220], [260], [261]; the training procedure for a WGAN with augmentations is

$$\mathcal{L}_D = \mathbb{E}_{\boldsymbol{x}\sim p_d(\boldsymbol{x})}[D(T(\boldsymbol{x}))] - \mathbb{E}_{\boldsymbol{z}\sim p(\boldsymbol{z})}[D(T(G(\boldsymbol{z})))], \tag{22a}$$

$$\mathcal{L}_G = \mathbb{E}_{\boldsymbol{z}\sim p(\boldsymbol{z})}[D(T(G(\boldsymbol{z})))], \tag{22b}$$

where T is a random augmentation. These approaches have been shown to improve sample quality on equivalent architectures and stabilise training. Each work offers a different perspective on why augmentation is so effective: the increased quantity of training data in conjunction with the more difficult discrimination task prevents overfitting and in turn collapse [18], notably this applies even on very small datasets (100 samples); the nature of GAN training leads to the generated and data distributions having non-overlapping supports, complicating training [208], strong augmentations may cause these distributions to overlap further. If an augmentation is differentiable and represents an invertible transformation of the data space's distribution, then the JS-divergence is invariant, and the generator is guaranteed to not create augmented samples [114], [220].

4.1.4 Discriminator Driven Sampling

In order to improve sample quality and address overpowered discriminators, numerous works have taken inspiration from the connection between GANs and energy models [256]. Interpreting the discriminator of a Wasserstein GAN [5] as an energy-based model means samples from the generator can be used to initialise an MCMC sampling chain which converges to the density learned by the discriminator, correcting errors learned by the generator [159], [223]. This is similar to pure EBM approaches, however, training the two networks adversarially changes the dynamics. The slow convergence rates of high dimensional MCMC sampling has led others to instead sample in the latent space [23], [205].

4.1.5 GANs Without Competition

Originally proposed as a proxy to measure GAN convergence [68], the duality gap is an upper bound on the JS-divergence that can be directly optimised [67], defined as

$$DG(D,G) = \max_{D'} V(G,D') - \min_{G'} V(G',D). \tag{23}$$

Cooperative training simplifies the optimisation procedure, avoiding oscillations. Each training step, however, requires optimising for D' and G' which slows down training and could suffer from vanishing gradients.

4.2 Architectures

Careful network design is a key component for stable GAN training. Scaling any deep neural network to high-resolution data is non-trivial due to vanishing gradients and high memory usage, but since the discriminator can classify high-resolution data more easily, GANs notably struggle [170].

Early approaches designed hierarchical architectures, dividing the learning procedure into more easily learnable chunks. LapGAN [39] builds a Laplacian pyramid such that at each layer, a GAN conditioned on the previous image resolution predicts a residual adding detail. Stacked GANs [98], [254] use two GANs trained successively: the first generates low-resolution samples, then the second upsamples and corrects the first, thus fewer GANs need to be trained. A related approach, progressive growing [113], [116], iteratively trains a single GAN at higher resolutions by adding layers to both the generator and discriminator upscaling the previous output, after the previous resolution converges. Training in this manner, however, not only takes a long time but leads to high frequency components being learned in the lower layers, resulting in shift artefacts [117].

Accordingly, a number of works have targeted a single GAN that can be trained end-to-end. DCGAN [180] introduced a fully convolutional architecture with batch normalisation [101] and ReLU/LeakyReLU activations. BigGAN [18] employ a number of tricks to scale to high resolutions including using very large mini-batches to reduce variation, spectral normalisation to discourage spectral collapse, and using large datasets to prevent overfitting. Despite this, training collapse still occurs thus requiring early stopping. Another approach is to include skip connections between the generator and discriminator at each resolution, allowing gradients to flow through shorter paths to each layer, providing extra information to the generator [112], [117], [228]. By treating subsets of the generator's parameters as smaller generators, Anycost GANs extend this approach, allowing samples to be generated at multiple resolutions and speeds[136]. To learn long-range dependencies, GANs can be built with self-attention components [104], [234], [253], however, full quadratic attention does not scale well to high dimensional data.

4.3 Training Speed

The mini-max nature of GAN training leads to slow convergence, if achieved at all. This problem has been exacerbated by numerous works as a byproduct of improving stability or sample quality. One such example is that by using very large mini-batches, reducing variance and covering more modes, sample quality can be improved significantly, however, this comes at the cost of slower training [18]. Small-GAN [195] combats this by replacing large batches with small batches that approximate the shape of the larger batch using core set sampling [195], significantly improving the mode coverage and sample quality of GANs trained with small batches.

While strong discriminator regularisation stabilises training, it allows the generator to make small changes and trick the discriminator, making convergence very slow. Rob-GAN [140], include an adversarial attack step [148] that perturbs real images to trick the discriminator without altering the content inordinately, adapting the GAN objective into a min-max-min problem. This provides a weaker regularisation,

enforcing small Lipschitz values locally rather than globally. This approach has been connected with the follow-the-ridge algorithm [240], [262], an optimisation approach for solving mini-max problems that reduces the optimisation path and converges to local mini-max points.

Another approach to improve training speed is to design more efficient architectures. Depthwise convolutions [32] apply separate convolutions to each channel of a tensor reducing the number of operations and hence also the runtime, have been found to have comparable quality to standard convolutions [160]. Lightweight GANs [137] achieve fast training using a number of tricks including small batch sizes, skip-layer excitation modules which provide efficient shortcut gradient flow, as well as using a self-supervised discriminator forcing good features to be learned.

5 Autoregressive Likelihood Models

Autoregressive generative models [9] are based on the chain rule of probability, where the probability of a variable that can be decomposed as $\boldsymbol{x} = x_1, \ldots, x_n$ is expressed as

$$p(\boldsymbol{x}) = p(x_1, \ldots, x_n) = \prod_{i=1}^{n} p(x_i | x_1, \ldots, x_{i-1}). \tag{24}$$

As such, unlike GANs and energy models, it is possible to directly maximise the likelihood of the data by training a recurrent neural network to model $p(x_i|\boldsymbol{x}_{1:i-1})$ by minimising the negative log-likelihood,

$$-\ln p(\boldsymbol{x}) = -\sum_{i}^{n} \ln p(x_i | x_1, \ldots, x_{i-1}). \tag{25}$$

While autoregressive models are extremely powerful density estimators, sampling is inherently a sequential process and can be exceedingly slow on high dimensional data. Additionally, data must be decomposed into a fixed ordering; while the choice of ordering can be clear for some modalities (e.g. text and audio), it is not obvious for others such as images and can affect performance depending on the network architecture used.

5.1 Architectures

The majority of research is focused on improving network architectures to increase their receptive fields and memory, ensuring the network has access to all parts of the input to encourage consistency, as well as increasing the network capacity, allowing more complex distributions to be modelled.

5.1.1 Masked Multilayer Perceptrons

One approach to build autoregressive models is to mask the weights of simple multilayer perceptron (MLP) autoencoders so as to satisfy the autoregressive property. The neural autoregressive density estimator (NADE) [130], which can be viewed as a mean-field approximation of a restricted Boltzmann machine, achieves this for binary data by placing time-dependent masks on an MLP with one hidden layer. Specifically, at time step i, weights are masked so that the entire hidden state $\boldsymbol{h}_i$ and output $p(x_i|\boldsymbol{x}_{<i})$ are dependent only on $\boldsymbol{x}_{<i}$; formally this can be defined as

$$p(x_i = 1|\boldsymbol{x}_{<i}) = \sigma(b_i + (\boldsymbol{W}^T)_{i,:}\boldsymbol{h}_i), \tag{26a}$$

$$\boldsymbol{h}_i = \sigma(\boldsymbol{c} + \boldsymbol{W}_{:,<i}\boldsymbol{x}_{<i}), \tag{26b}$$

where $\boldsymbol{W}_{:,<d}$ is the first $d-1$ columns of a shared weight matrix $\boldsymbol{W}$, and b_i and $\boldsymbol{c}$ are biases. The RNADE [225] generalises NADE to real valued data by instead modelling $p(x_i|\boldsymbol{x}_{<i})$ with mixture distributions parameterised by the network. An alternative masking procedure known as MADE [56] allows for parallel density estimation by placing a mask fixed over time on an MLP so that no connections exist between $p(x_i|\boldsymbol{x}_{<i})$ and $\boldsymbol{x}_{\geq i}$. Additionally, MADE is more readily vectorisable and does not suffer from neuron saturation since the number of inputs to all neurons is constant with respect to time.

5.1.2 Recurrent Neural Networks

A natural architecture to apply is that of standard recurrent neural networks (RNNs) such as LSTMs [87], [212], [232] and GRUs [34], [151] which model sequential data by tracking information in a hidden state. However, RNNs are known to forget information, limiting their receptive field thus preventing modelling of long range relationships. This can be improved by stacking RNNs that run at different frequencies allowing long data such as multiple seconds of audio to be modelled [34]. Nevertheless, their sequential nature means that training can be too slow for many tasks.

5.1.3 Causal Convolutions

An alternative approach is that of causal convolutions, which apply masked or shifted convolutions over a sequence [29], [188], [231]. When stacked, this only provides a receptive field linear with depth, however, by dilating the convolutions to skip values with some step the receptive field can be orders of magnitude higher.

5.1.4 Self-Attention

Neural attention is an approach which at each successive time step is able to select where it wishes to 'look' at previous time steps. This concept has been used to autoregressively 'draw' images onto a blank 'canvas' [66] in a manner similar to human drawing. More recently self-attention (known as Transformers when used in an encoder-decoder setup) [234] has made significant strides improving not only autoregressive models, but also other generative models due to its parallel nature, stable training, and ability to effectively learn long-distance dependencies. This is achieved using an attention scheme that can reference any previous input where an entirely independent process is used per time step so that there are no dependencies. Specifically, inputs are encoded as key-value pairs, where the values $\boldsymbol{V}$ represent the inputs, and the keys $\boldsymbol{K}$ act as an indexing method. At each time step a query q is made; taking the dot product of the queries and keys, a similarity vector is formed that describes which value vectors to access. This process can be expressed as

$$\text{Attention}(\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V}) = \text{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{d_k}}\right)\boldsymbol{V}, \tag{27}$$

where d_k is the key/query dimension and is used to normalise gradient magnitudes. Since the self-attention process

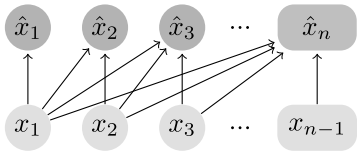

Fig. 6. Autoregressive models decompose data points using the chain rule and learn conditional probabilities.

contains no recurrence, positional information must be passed into the function. A simple effective method to achieve this is to add sinusoidal positional encodings which combine sine and cosine functions of different frequencies to encode positional information [234]; alternatively others use trainable positional embeddings [31].

The infinite receptive fields of attention provides a powerful tool for representing data, however, the attention matrix $\boldsymbol{Q}\boldsymbol{K}^T$ grows quadratically with data dimension, making scaling difficult. Approaches include scaling across large quantities of GPUs [19], interleaving attention between causal convolutions [29], attending over local regions [177], and using sparse attention patterns that provide global attention when multiple layers are stacked [31]. More recently, a number of linear transformers have been proposed whose memory and time footprints grow linearly with data dimension [33], [118], [239]. By approximating the softmax operation with a kernel function with feature representation $\phi(\boldsymbol{x})$, the order of multiplications can be rearranged to

$$\big(\phi(\boldsymbol{Q})\phi(\boldsymbol{K})^T\big)\boldsymbol{V} = \phi(\boldsymbol{Q})\big(\phi(\boldsymbol{K})^T\boldsymbol{V}\big), \tag{28}$$

allowing $\phi(\boldsymbol{K})^T\boldsymbol{V}$ to be cached and used for each query.

5.1.5 Multiscale Architectures

Even with a linear autoregressive model, $O(N)$ for N pixels, scaling to high-resolution images grows quadratically with resolution. One multi-scale approach reduces this complexity to $O(\ln N)$ by successively upscaling images, making the assumption that when upscaling, each pixel is dependent only on its adjacent area and the previous resolution image, allowing scaling to high resolutions [183]. To avoid making independence assumptions, [154] partition images in an interleaving pattern so that sub-images are the same size and capture global structure. Sub-images are generated autoregressively pixel-wise and are conditioned on previously generated sub-images; while this reduces the memory required, sampling times are still slow.

5.2 Data Modelling Decisions

When generating text, output variables are often modelled using a multinomial distribution since tokens are discrete and are in general unrelated. However, this modelling assumption can cause complications or be infeasible in other cases such as 16-bit audio modelling, in which magnitude would not be intrinsically modelled and 65,536 output neurons would be required. Solutions proposed include:

- Applying μ-law, a logarithmic companding algorithm which takes advantage of human perception of sound, then quantizing to 8-bit values [229].
- First predicting the first 8-bits, then predicting the second 8-bits conditioned on the first.

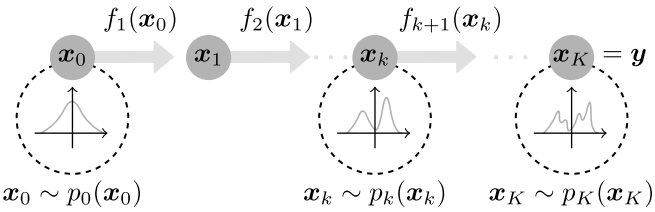

Fig. 7. Normalizing flows build complex distributions by mapping a simple distribution through invertible functions.

- Modelling output probabilities using a mixture of logistic distributions (MoL) has the benefits of providing more useful gradients and allowing intensities never seen to still be sampled [188].

Nevertheless, these assumptions restrict the expressiveness of the network, for instance, MoLs struggle to model high frequency signals as found in raw image data; a simple solution in this case is to add Gaussian noise, reducing the Lipschitz constant of the data distribution [152]. This restriction can be removed at the expense of less efficient sampling by learning an autoregressive energy model, for instance, by approximating normalising constants [158] or through score matching [153]. Alternatively, quantile regression, which minimises Wasserstein distance, can be used to learn an approximation of the inverse cumulative distribution [172].

When modelling images, many works use "raster scan" ordering [188], [231], [232] where pixels are estimated row by row. Alternatives have been proposed such as "zig-zag" ordering [29] which allows pixels to depend on previously sampled pixels to the left and above, providing more relevant context. Another factor when modelling images is how to factorise sub-pixels. While it is possible to treat them as independent variables, this adds additional complexity. Alternatively, it is possible to instead condition on whole pixels, and output joint distributions in a single step [188].

6 Normalizing Flows

While training autoregressive models through maximum likelihood offers plenty of benefits including stable training, density estimation, and a useful validation metric, the slow sampling speed and poor scaling properties handicaps them significantly. Normalizing flows are a technique that also allows exact likelihood calculation while being efficiently parallelisable as well as offering a useful latent space for downstream tasks. Consider an invertible, smooth function $f : \mathbb{R}^d \to \mathbb{R}^d$; by applying this transformation to a random variable $\boldsymbol{x} \sim p(\boldsymbol{x})$, then the distribution of the resulting random variable $\boldsymbol{y} = f(\boldsymbol{x})$ can be determined through the change of variables rule (and application of the chain rule),

$$p(\boldsymbol{y}) = p(\boldsymbol{x})\left|\det\frac{\partial f^{-1}}{\partial \boldsymbol{y}}\right| = p(\boldsymbol{x})\left|\det\frac{\partial f}{\partial \boldsymbol{x}}\right|^{-1}. \tag{29}$$

Consequently, arbitrarily complex densities can be constructed by composing simple maps and applying Eq. (29) [235]. This chain is known as a normalizing flow [184] (see Fig. 7). The density $p_K(\boldsymbol{x}_K)$ obtained by successively transforming a random variable $\boldsymbol{x}_0$ with distribution p_0 through a chain of K transformations f_k can be defined as

TABLE 3
Normalizing Flow Layers: $\odot$ Represents Elementwise Multiplication, $\star_l$ Represents a Cross-Correlation Layer

Description	Function	Inverse Function	Log-Determinant
Low Rank			
Planar [185]	$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{u}h(\boldsymbol{w}^T\boldsymbol{z} + b)$ With $\boldsymbol{w} \in \mathbb{R}^D$, $\boldsymbol{u} \in \mathbb{R}^D$, $\boldsymbol{b} \in \mathbb{R}$	No closed form inverse	$\ln\lvert 1 + \boldsymbol{u}^T h'(\boldsymbol{w}^T\boldsymbol{z} + b)\boldsymbol{w}\rvert$
Sylvester [14], [79]	$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{U}h(\boldsymbol{W}^T\boldsymbol{x} + \boldsymbol{b})$	No closed form inverse	$\ln\det(\boldsymbol{I}_M + \mathrm{diag}(h'(\boldsymbol{W}^T\boldsymbol{x} + \boldsymbol{b}))\boldsymbol{W}\boldsymbol{U}^T)$
Coupling/Autoregressive			
General Coupling	$\boldsymbol{y}^{(1:d)} = \boldsymbol{x}^{(1:d)}$ $\boldsymbol{y}^{(d+1:D)} = h(\boldsymbol{x}^{(d+1:D)}; f_\theta(\boldsymbol{x}^{(1:d)}))$	$\boldsymbol{x}^{(1:d)} = \boldsymbol{y}^{(1:d)}$ $\boldsymbol{x}^{(d+1:D)} = h^{-1}(\boldsymbol{y}^{(d+1:D)}; f_\theta(\boldsymbol{y}^{(1:d)}))$	$\ln\lvert\det \nabla_{\boldsymbol{x}^{(d+1:D)}} h\rvert$
MAF [177]	$y^{(t)} = h(x^{(t)}; f_\theta(\boldsymbol{x}^{(1:t-1)}))$	$x^{(t)} = h^{-1}(y^{(t)}; f_\theta(\boldsymbol{x}^{(1:t-1)}))$	$-\sum_{t=1}^{D} \ln\lvert\frac{\partial y^{(t)}}{\partial x^{(t)}}\rvert$
IAF [125]	$y^{(t)} = h(x^{(t)}; f_\theta(\boldsymbol{y}^{(1:t-1)}))$	$x^{(t)} = h^{-1}(y^{(t)}; f_\theta(\boldsymbol{y}^{(1:t-1)}))$	$\sum_{t=1}^{D} \ln\lvert\frac{\partial y^{(t)}}{\partial x^{(t)}}\rvert$
Affine Coupling [43]	$h(\boldsymbol{x}; \boldsymbol{\theta}) = \boldsymbol{x} \odot \exp(\boldsymbol{\theta}_1) + \boldsymbol{\theta}_2$	$h^{-1}(\boldsymbol{y}; \boldsymbol{\theta}) = (\boldsymbol{y} - \boldsymbol{\theta}_2) \odot \exp(-\boldsymbol{\theta}_1)$	$\sum_{i=1}^{d} \theta_1^{(i)}$
Flow++ [86]	$h(\boldsymbol{x}; \boldsymbol{\theta}) = \exp(\boldsymbol{\theta}_1) \odot F(\boldsymbol{x}, \boldsymbol{\theta}_3) + \boldsymbol{\theta}_2$ where F is a monotone function.	Calculated through bisection search	$\sum_{i=1}^{d} \theta_1^{(i)} + \ln\frac{\partial F(\boldsymbol{x}, \boldsymbol{\theta}_3)_i}{\partial x_i}$
Spline Flows [158][50][51]	$h(\boldsymbol{x}; \boldsymbol{\theta}) = \mathrm{Spline}(\boldsymbol{x}; \boldsymbol{\theta})$ where $\boldsymbol{\theta}$ are the spline's knots.	$h^{-1}(\boldsymbol{y}; \boldsymbol{\theta}) = \mathrm{Spline}^{-1}(\boldsymbol{y}; \boldsymbol{\theta})$	Computed in closed-form as a product of quotient derivatives
B-NAF [39]	$\boldsymbol{y} = \boldsymbol{W}\boldsymbol{x}^T$ for blocked weights: $\boldsymbol{W} = \exp(\tilde{\boldsymbol{W}}) \odot \boldsymbol{M}_d + \tilde{\boldsymbol{W}} \odot \boldsymbol{M}_o$ where $\boldsymbol{M}_d$ selects diagonal blocks and $\boldsymbol{M}_o$ selects off-diagonal blocks.	No closed form inverse	$\ln\sum_{i=1}^{d} \exp(\tilde{W}_{ii})$
Convolutions			
1x1 Convolution [124]	$h \times w \times c$ tensor $\boldsymbol{x}$ & $c \times c$ tensor $\boldsymbol{W}$ $\forall i, j : \boldsymbol{y}_{i,j} = \boldsymbol{W}\boldsymbol{x}_{i,j}$	$\forall i, j : \boldsymbol{x}_{i,j} = \boldsymbol{W}^{-1}\boldsymbol{y}_{i,j}$	$h \cdot w \cdot \ln\lvert\det \boldsymbol{W}\rvert$
Emerging Convolutions [91]	$\boldsymbol{k} = \boldsymbol{w}_1 \odot \boldsymbol{m}_1, \quad \boldsymbol{g} = \boldsymbol{w}_2 \odot \boldsymbol{m}_2$ $\boldsymbol{y} = \boldsymbol{k} \star_l (\boldsymbol{g} \star_l \boldsymbol{x})$	$\boldsymbol{z_t} = (\boldsymbol{y}_t - \sum_{i=t+1} G_{t,i}\boldsymbol{z_i})/G_{t,t}$ $\boldsymbol{x}_t = (\boldsymbol{z_t} - \sum_{i=1}^{t-1} K_{t,i}\boldsymbol{x}_i)/K_{t,t}$	$\sum_c \ln\lvert\boldsymbol{k}_{c,c,m_y,m_x}\boldsymbol{g}_{c,c,m_y,m_x}\rvert$
Lipshitz Residual			
i-ResNet [8]	$\boldsymbol{y} = \boldsymbol{x} + f(\boldsymbol{x})$ where $\lVert f\rVert_L < 1$	$\boldsymbol{x}_1 = \boldsymbol{y}. \quad \boldsymbol{x}_{n+1} = \boldsymbol{y} - f(\boldsymbol{x}_n)$ converging at an exponential rate	$\mathrm{tr}(\ln(\boldsymbol{I} + \nabla_{\boldsymbol{x}} f)) =$ $\sum_{k=1}^{\infty} (-1)^{k+1} \frac{\mathrm{tr}((\nabla_{\boldsymbol{x}} f)^k)}{k}$

$$\boldsymbol{x}_K = f_K \circ \cdots \circ f_2 \circ f_1(\boldsymbol{x}_0), \tag{30a}$$

$$\ln p_K(\boldsymbol{x}_K) = \ln p_0(\boldsymbol{x}_0) - \sum_{k=1}^{K} \ln \left| \det \frac{\partial f_k}{\partial \boldsymbol{x}_{k-1}} \right|. \tag{30b}$$

Each transformation therefore must be sufficiently expressive while being easily invertible and have an efficient to compute Jacobian determinant. While restrictive, there have been a number of works which have introduced more powerful invertible functions (see Table 3). Nevertheless, normalizing flow models are typically less parameter efficient than other generative models.

One disadvantage of requiring transformations to be invertible is that the input dimension must be equal to the output dimension which makes deep models inefficient and difficult to train. A popular solution to this is to use a multi-scale architecture [42], [123] (see Fig. 8) which divides the process into a number of stages, at the end of each half of the remaining units are factored out and treated immediately as outputs. This allows latent variables to sequentially represent course to fine features and permits deeper architectures.

6.1 Coupling and Autoregressive Layers

A simple way of building an expressive invertible function is the coupling flow [41], which divide inputs into two and applies a bijection h on one half parameterised by the other,

$$\boldsymbol{y}^{(1:d)} = \boldsymbol{x}^{(1:d)}, \tag{31a}$$

$$\boldsymbol{y}^{(d+1:D)} = h(\boldsymbol{x}^{(d+1:D)}; f_\theta(\boldsymbol{x}^{(1:d)})), \tag{31b}$$

here f can be arbitrarily complex i.e. a neural network. h tends to be selected as an elementwise function making the Jacobian triangular allowing efficient computation of the determinant, i.e. the product of elements on the diagonal.

6.1.1 Affine Coupling

A simple example of this is the affine coupling layer [42],

$$\boldsymbol{y}^{(d+1:D)} = \boldsymbol{x}^{(d+1:D)} \odot \exp(f_\sigma(\boldsymbol{x}^{(1:d)})) + f_\mu(\boldsymbol{x}^{(1:d)}), \tag{32}$$

which has a simple Jacobian determinant and can be trivially rearranged to obtain a definition of $\boldsymbol{x}^{(d+1:D)}$ in terms of

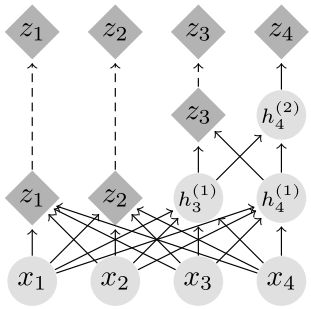

Fig. 8. Factoring out variables at different scales allows normalizing flows to scale to high dimensional data.

$\boldsymbol{y}$, provided that the scaling coefficients are not 0. This simplicity, however, comes at the cost of expressivity; while stacking numerous such flows increases their expressivity, allowing them to learn representations of complex high dimensional data such as images [123], it is unknown whether multiple affine flows are universal approximators [175].

6.1.2 Monotone Functions

Another method of creating invertible functions that can be applied element-wise is to enforce monotonicity. One possibility to achieve this is to define h as an integral over a positive but otherwise unconstrained function g [242],

$$h(x_i;\boldsymbol{\theta}) = \int_0^{x_i} g_{\boldsymbol{\phi}}(x;\boldsymbol{\theta}_1)dx + \theta_2, \tag{33}$$

however, this integration requires numerical approximation. Alternatively, by choosing g to be a function with a known integral solution, h can be efficiently evaluated. This has been accomplished using positive polynomials [102] and the CDF of a mixture of logits [85]. Both cases, however, don't have analytical inverses and have to be approximated iteratively with bisection search. Another option is to represent g as a monotonic spline: a piecewise function where each piece is easy to invert. As such, the inverse is as fast to evaluate as the forward pass. Linear and quadratic splines [157], cubic splines [49], and rational-quadratic splines [50] have been applied so far.

6.1.3 Autoregressive Flows

For a single coupling layer, a significant proportional of inputs remain unchanged. A more flexible generalisation of coupling layers is the autoregressive flow, or MAF [176],

$$y^{(t)} = h(x^{(t)}; f_\theta(\boldsymbol{x}^{(1:t-1)})). \tag{34}$$

Here f_θ can be arbitrarily complex, allowing the use of advances in autoregressive modelling (Section. 5), and h is a bijection as used for coupling layers. Some monotonic bijectors have been created specifically for autoregressive flows, namely Neural Autoregressive Flows (NAF) [95] and Block NAF [38]. Unlike coupling layers, a single autoregressive flow is a universal approximator.

Alternatively, an autoregressive flow can be conditioned on $\boldsymbol{y}^{(1:t-1)}$ rather than $\boldsymbol{x}^{(1:t-1)}$, this is known as an Inverse Autoregressive Flow, or IAF [124]. While coupling layers can be evaluated efficiently in both directions, MAF permits parallel density estimation but sequential sampling, and IAF permits parallel sampling but sequential density estimation.

6.1.4 Probability Density Distillation

Inverse autoregressive flows [124] offer the ability to sample from an autoregressive model in parallel, however, training via maximum likelihood is inherently sequential making this infeasible for high dimensional data. Probability density distillation [230] has been proposed as a solution to this where a second pre-trained autoregressive network is used as a 'teacher' network while an IAF network is used as a 'student' and mimics the teacher's distribution by minimising the KL divergence between the two distributions

$$D_{KL}(p_S||p_T) = H(p_S, p_T) - H(p_S), \tag{35}$$

where p_S and p_T are the student's and teacher's distributions respectively, $H(p_S, p_T)$ is the cross-entropy between p_S and p_T, and $H(p_S)$ is the entropy of p_S. Crucially, this never requires the student's inverse function to be used allowing it to be computed entirely in parallel.

6.2 Convolutional

A considerable problem with coupling and autoregressive flows is the restricted triangular Jacobian, meaning that all inputs cannot interact with each other. Simple solutions involve fixed permutations on the output space such as reversing the order [41], [42]. A more general approach is to use a 1×1 convolution which is equivalent to a linear transformation applied across channels [123]. Numerous works have been proposed to generalise these to larger kernel sizes. A number of these apply variations on causal convolutions [229], including emerging convolutions [90] whose inverse is sequential, MaCow [144] which uses smaller conditional fields allowing more efficient sampling, and MintNet [203] which approximates the inverse using fixed-point iteration. Alternative approaches to causal masking involve imposing repeated (periodic) structure [111], however in general this is not a good assumption for image modelling, as well as representing convolutions as exponential matrix-vector products, $\exp(M)\boldsymbol{x}$, approximated implicitly with a power series, allowing otherwise unconstrained kernels [91].

6.3 Residual Flows

Residual networks [79] are a popular technique to build deep neural networks that alleviate the vanishing gradients problem. By restricting f_θ, invertible residual networks can be built by stacking blocks of the form

$$\boldsymbol{y} = \boldsymbol{x} + f_\theta(\boldsymbol{x}). \tag{36}$$

6.3.1 Matrix Determinant Lemma

If a function has a certain residual form, then its Jacobian determinant can be computed with the matrix determinant lemma [184]. A simple example is planar flow [184] which is equivalent to a 3 layer MLP with a single neuron bottleneck

$$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{u}h(\boldsymbol{w}^T\boldsymbol{x} + b), \tag{37}$$

where $\boldsymbol{u}, \boldsymbol{w} \in \mathbb{R}^d$, $b \in \mathbb{R}$, and h is a differentiable non-linearity function. Planar flows are invertible provided some simple conditions are satisfied, however its inverse is difficult to compute making it only practical for density estimation tasks. A higher rank generalisation of the matrix determinant lemma has been applied to planar flows, known as Sylvester flows, removing the severe bottleneck thus allowing greater representation ability [13], [78].

6.3.2 Lipschitz Constrained

By restricting the Lipschitz constant of f_θ, $\|f_\theta\|_L < 1$, then this block is invertible [8]. The inverse, however, has no closed form definition but can be found through fixed-point iteration which by the Banach fixed-point theorem converges to a fixed unique solution at an exponential rate dependant on $\|f_\theta\|_L$. The authors originally proposed a biased approximation of the log determinant of the Jacobian as a power series where the Jacobian trace is approximated using Hutchkinson's trace estimator (see Table 3), but an unbiased approximator known as a Russian roulette estimator has also been proposed [25]. Unlike coupling layers, residual flows have dense Jacobians, allowing interaction. Enforcing Lipschitz constraints has been achieved with convolutional networks [59], [138], [156] as well as self-attention [119].

Making strong Lipschitz assumptions severely restricts the class of functions learned; an N layer residual flow network is at most 2^N-Lipshitz. Implicit flows [142] bypass this by solving implicit equations of the form

$$F(\boldsymbol{x}, \boldsymbol{y}) = f_\theta(\boldsymbol{x}) - f_\phi(\boldsymbol{y}) + \boldsymbol{x} - \boldsymbol{y} = \boldsymbol{0}, \tag{38}$$

where both f_θ and f_ϕ both have Lipschitz constants less than 1. Both the forwards (solve for $\boldsymbol{y}$ given $\boldsymbol{x}$) and backwards (solve for $\boldsymbol{x}$ given $\boldsymbol{y}$) directions require solving a root finding problem similar to the inverse process of residual flows; indeed, an implicit flow is equivalent to the composition of a residual flow and the inverse of a residual flow. This allows them to model arbitrary Lipschitz transformations.

6.4 Surjective and Stochastic Layers

Restricting the class of functions available to those that are invertible introduces a number of practical problems related to the topology-preserving property of diffeomorphisms. For example, mapping a uni-modal distribution to a multi-modal distribution is extremely challenging, requiring a highly varying Jacobian [43]. By composing bijections with surjective or stochastic layers these topological constraints can be bypassed [163]. While the log-likelihood of stochastic layers can only be bounded by their ELBO, functions surjective in the inference direction permit exact likelihood evaluation even with altered dimensionality. Surjective transformations have the following likelihood contributions

$$\mathbb{E}_{q(\boldsymbol{y}|\boldsymbol{x})}\left[\ln \frac{p(\boldsymbol{x}|\boldsymbol{y})}{q(\boldsymbol{y}|\boldsymbol{x})}\right], \tag{39}$$

where $p(\boldsymbol{x}|\boldsymbol{y})$ is deterministic for generative surjections, and $q(\boldsymbol{y}|\boldsymbol{x})$ is deterministic for inference surjections.

One approach to build a surjective layer is to augment the input space with additional dimensions allowing smoother transformation to be learned [24], [47], [94]; the inverse process, where some dimensions are factored out, is equivalent to a multi-scale architecture [42]. Another approach known as RAD [43] learns a partitioning of the data space into disjoint subsets $\{\mathcal{Y}_i\}_{i=1}^K$, and applies piecewise bijections to each region $g_i : \mathcal{X} \to \mathcal{Y}_i, \forall i \in \{1, \ldots, K\}$. The generative direction learns a classifier on $\mathcal{X}$, $i \sim p(i|\boldsymbol{x})$, allowing the inverse to be calculated as $\boldsymbol{y} = g_i(\boldsymbol{x})$. Similar to both of these approaches are CIFs [35] which consider a continuous partitioning of the data space via augmentation equivalent to an infinite mixture of normalizing flows. Other approaches include modelling finite mixtures of flows [46].

Some powerful stochastic layers have already been discussed in this survey, namely VAEs [122] and DDPMs [86]. Stochastic layers have been incorporated into normalizing flows by interleaving small energy models, sampled with MCMC, between bijectors [244].

6.5 Discrete Flows

The normalizing flow framework can be extended to discrete distributions, by restricting transformation functions to be discrete e.g. $f : \mathcal{X}^d \to \mathcal{X}^d$. Integer discrete flows (IDF) achieve this using additive coupling layers, rounding translation values to the nearest integer and approximating gradients with the straight-through estimator [93]; discrete flows [219] apply affine coupling layers in modulo space while also restricting the translation and scaling coefficients to a finite number of possible values. In this case the change of variables rule (Eq. (29)) simplifies to [93], [219]

$$p(\boldsymbol{x}) = p(f(\boldsymbol{x})). \tag{40}$$

Unlike the continuous case, there is no Jacobian determinant term; intuitively this term adjusts for volume changes, however, in a discrete space there is no volume. As such, there is no requirement for f to have an efficiently computable Jacobian determinant [219]. The absence of this term is restricting, however, discrete flows can only permute the values of $p(\boldsymbol{x})$, not change them i.e. a uniform base distribution can only be mapped to another uniform distribution [175]. Nevertheless, this can be avoided by embedding the data into a space with more values than the data, making IDFs more flexible than discrete flows [12].

6.6 Continuous Time Flows

It is possible to consider a normalizing flow with an infinite number of steps that is defined instead by an ordinary differential equation specified by a Lipschitz continuous neural network f with parameters θ, that describes the transformation of a hidden state $\boldsymbol{x}(t) \in \mathbb{R}^D$ [26],

$$\frac{\partial \boldsymbol{x}(t)}{\partial t} = f(\boldsymbol{x}(t), t, \theta). \tag{41}$$

Starting from input noise $\boldsymbol{x}(t_0)$, an ODE solver can solve an initial value problem for some time t_1, at which data is

defined, $\boldsymbol{x}(t_1)$. Modelling a transformation in this form has a number of advantages such as inherent invertibility by running the ODE solver backwards, parameter efficiency, and adaptive computation. However, it is not immediately clear how to train such a model through backpropagation. While it is possible to backpropagate directly through an ODE solver, this limits the choice of solvers to differentiable ones as well as requiring large amounts of memory. Instead, the authors apply the adjoint sensitivity method which instead solves a second, augmented ODE backwards in time and allows the use of a black box ODE solver. That is, to optimise a loss dependent on an ODE solver

$$\begin{aligned}\mathcal{L}(\boldsymbol{x}(t_1)) &= \mathcal{L}\bigg(\boldsymbol{x}(t_0) + \int_{t_0}^{t_1} f(\boldsymbol{z}(t), t, \theta)dt\bigg), \\ &= \mathcal{L}(\text{ODESolve}(\boldsymbol{x}(t_0), f, t_0, t_1, \theta)),\end{aligned} \tag{42}$$

the adjoint $\boldsymbol{a}(t) = \frac{\partial \mathcal{L}}{\partial \boldsymbol{x}(t)}$ can be used to calculate the derivative of loss with respect to the parameters in the form of another initial value problem [178],

$$\frac{\partial \mathcal{L}}{\partial \theta} = \int_{t_1}^{t_0} \Big(\frac{\partial \mathcal{L}}{\partial \boldsymbol{x}(t)}\Big)^T \frac{\partial f(\boldsymbol{x}(t), t, \theta)}{\partial \theta} dt, \tag{43}$$

which can be efficiently evaluated by automatic differentiation at a time cost similar to evaluating f itself.

Despite the complexity of this transformation, the continuous change of variables rule is remarkably simple

$$\frac{\partial \ln p(\boldsymbol{x}(t))}{\partial t} = -tr\bigg(\frac{\partial}{\partial \boldsymbol{x}(t)} f(\boldsymbol{x}(t), t, \theta)\bigg), \tag{44}$$

and can be computed using an ODE solver as well. The resulting continuous-time flow is known as FFJORD [61]. Since the length of the flow tends to infinity (an infinitesimal flow), the true posterior distribution can be recovered [184].

As previously mentioned, invertible functions suffer from topological problems; this is especially true for Neural ODEs since their continuous nature prevents trajectories from crossing. Similar to augmented normalizing flows [94], this can be solved by providing additional dimensions for the flow to traverse [47]. Specifically, a p-dimensional Euclidean space can be approximated by a Neural ODE in a $(2p+1)$-dimensional space [252].

6.6.1 Regularising Trajectories

ODE solvers can require large numbers of network evaluations, notably when the ODE is stiff or the dynamics change quickly in time. By introducing regularisation, a simpler ODE can be learned, reducing the number of evaluations required. Specifically, all works here are inspired by optimal transport theory to encourage straight trajectories. Monge-Ampère Flow [255] and Potential Flow Generators [248] parameterise a potential function satisfying the Monge-Ampère equation [21], [235] with a neural network. RNODE [53] applies transport costs to FFJORD as well as regularising the Frobenius norm of the Jacobian, encouraging straight trajectories. By combining these approaches, OT-Flow [171] utilises the optimal transport derivation to derive an exact trace definition with cost similar to stochastic estimators.

7 Evaluation Metrics

A huge problem when developing generative models is how to effectively evaluate and compare them. Qualitative comparison of random samples plays a large role in the majority of state-of-the-art works, however, it is subjective and time-consuming to compare many works. Calculating the log-likelihood on a separate validation set is popular for tractable likelihood models but comparison with implicit likelihood models is difficult and while it is a good measure of diversity, it does not correlate well with quality [213].

One approach to quantify sample quality is Inception Score (IS) [187] which takes a trained classifier and determines whether a sample has low label entropy, indicating that a meaningful class is likely, and whether the distribution of classes over a large number of samples has high entropy, indicating that a diverse range of images can be sampled. A perfect IS can be scored by a model that creates only one image per class [143] leading to the creation of Fréchet Inception Distance (FID) [80] which models the activations of a particular layer of a classifier as multivariate Gaussians for real and generated data, measuring the Fréchet distance between the two.

These approaches are trivially solved by memorising the dataset and are less applicable to non-natural image-related data. Kernel Inception Distance (KID) [14] instead calculates the squared maximum mean discrepancy in feature space, however, pretrained features may not be sufficient to detect overfitting. Another approach is to train a neural network to distinguish between real and generated samples similar to the discriminator from a GAN; while this detects overfitting, it increases the complexity and time required to evaluate a model and is biased towards adversarial models [73].

8 Applications

In general, the definition of a generative model means that any technique can be used on any modality/task, however, some models are more suited for certain tasks. Standard autoregressive networks are popular for text/audio generation [19], [31], [229]; VAEs have been applied but posterior collapse is difficult to mitigate [6], [17]; GANs are more parameter efficient but struggle to model discrete data [162] and suffer from mode collapse [127]; some normalizing flows offer parallel synthesis, providing substantial speedup [179], [219], [264]. Video synthesis is more challenging due its exceptionally high dimensionality, typically approaches combine a latent-based implicit generative model to generate individual frames, with an autoregressive network used to predict future latents [6], [128], [133] similar to how world models are constructed in reinforcement learning [75], [76]. Modality conversion has been achieved using GANs [263], VAE-GANs [139], and DDPMs [191].

8.1 Implicit Representation

Typically deep architectures discussed in this survey are built with data represented as discrete arrays thus using discrete components such as convolutions and self-attention. Implicit representation on the other hand treats data as

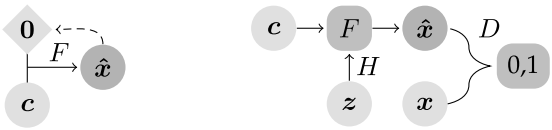

(a) Implicit GON [16]. (b) Implicit GAN [48].

Fig. 9. Implicit networks model data continuously permitting arbitrarily high resolutions. Dashed lines represent gradients, F is an implicit network, and H is a hypernetwork.

continuous signals, mapping coordinates to data values [196], [210]. Implicit Gradient Origin Networks (GONs; Fig. 9a) [16] form a latent variable model by concatenating latent vectors with coordinates which are passed through an implicit network; here latent vectors are calculated as the gradient of a reconstruction loss with respect to the origin. By sampling using a finer grid of coordinates, super-resolution beyond resolutions seen during training is possible. Other approaches to learn an implicit generative model as a GAN include directly feeding latents through an implicit network with upsampling [115] and mapping latents to the weights of an implicit function using a hyper-network [48] (Fig. 9b).

9 Conclusion

While GANs have led the way in terms of sample quality for some time now, the gap between other approaches is shrinking; the diminished mode collapse and simpler training objectives make these models more enticing than ever, however, the number of parameters required in addition to slow run-times pose a substantial handicap. Despite this, recent work in hybrid models offers a balance between extremes at the expense of extra model complexity that hinders broader adoption. The varied connections between these systems mean that advances in one field inevitably benefit others, for instance, improved variational bounds are beneficial for VAEs, diffusion models, and surjective flows, and the application of innovative data augmentation strategies has been found to offer benefits across numerous model classes without necessitating more powerful architectures. When it comes to scaling models to high-dimensional data, attention is a common theme, allowing long-range dependencies to be learned; recent advances in linear attention will aid scaling to even higher resolutions. Implicit networks are another promising direction, allowing efficient synthesis of arbitrarily high resolution and irregular data. Similar unified generative models capable of modelling continuous, irregular, and arbitrary length data, over different scales and domains will be key for the future of generalisation.

References

[1] G. Alain *et al.*, "GSNs: Generative stochastic networks," *Inf. Inference*, vol. 5, no. 2, pp. 210–249, 2016.

[2] A. Alemi, B. Poole, I. Fischer, J. Dillon, R. A. Saurous, and K. Murphy, "Fixing a broken ELBO," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 159–168.

[3] M. Arbel, L. Zhou, and A. Gretton, "Generalized energy based models," in *Proc. Int. Conf. Learn. Representations*, 2021.

[4] M. Arjovsky and L. Bottou, "Towards principled methods for training generative adversarial networks," in *Proc. Int. Conf. Learn. Representations*, 2017.

[5] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 214–223.

[6] M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine, "Stochastic variational video prediction," in *Proc. 6th Int. Conf. Learn. Representations*, 2018.

[7] M. Bauer and A. Mnih, "Resampled priors for variational autoencoders," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, 2019, pp. 66–75.

[8] J. Behrmann, W. Grathwohl, R. T. Q. Chen, D. Duvenaud, and J.-H. Jacobsen, "Invertible residual networks," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 573–582.

[9] Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin, "A neural probabilistic language model," *J. Mach. Learn. Res.*, vol. 3, pp. 1137–1155, 2003.

[10] Y. Bengio, N. Léonard, and A. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," 2013, *arXiv:1308.3432*.

[11] Y. Bengio, L. Yao, G. Alain, and P. Vincent, "Generalized denoising auto-encoders as generative models," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 899–907.

[12] R. van den Berg, A. A. Gritsenko, M. Dehghani, C. K. Sønderby, and T. Salimans, "IDF++: Analyzing and improving integer discrete flows for lossless compression," in *Proc. Int. Conf. Learn. Representations*, 2021.

[13] R. van den Berg, L. Hasenclever, J. M. Tomczak, and M. Welling, "Sylvester normalizing flows for variational inference," in *Proc. 34th Conf. Uncertainty Artif. Intell.*, 2018, pp. 393–402.

[14] M. Bińkowski, D. J. Sutherland, M. Arbel, and A. Gretton, "Demystifying MMD GANs," in *Proc. Int. Conf. Learn. Representations*, 2018.

[15] P. Bojanowski, A. Joulin, D. Lopez-Paz, and A. Szlam, "Optimizing the latent space of generative networks," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 600–609.

[16] S. Bond-Taylor and C. G. Willcocks, "Gradient origin networks," in *Proc. Int. Conf. Learn. Representations*, 2021.

[17] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio, "Generating sentences from a continuous space," 2016, *arXiv:1511.06349*.

[18] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," in *Proc. Int. Conf. Learn. Representations*, 2018.

[19] T. B. Brown *et al.*, "Language models are few-shot learners," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 1877–1901.

[20] Y. Burda, R. B. Grosse, and R. Salakhutdinov, "Importance weighted autoencoders," in *Proc. Int. Conf. Learn. Representations*, 2016.

[21] L. A. Caffarelli and M. Milman, *Monge Ampère Equation: Applications to Geometry, Optimization*. Providence, RI, USA: Amer. Math. Soc., 1999.

[22] M. A. Carreira-Perpinan and G. E. Hinton, "On contrastive divergence learning," in *Proc. 10th Int. Workshop Artif. Intell. Statist.*, 2005, pp. 33–40.

[23] T. Che *et al.*, "Your GAN is secretly an energy-based model and you should use discriminator driven latent sampling," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020.

[24] J. Chen, C. Lu, B. Chenli, J. Zhu, and T. Tian, "VFlow: More expressive generative flows with variational data augmentation," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 1660–1669.

[25] R. T. Q. Chen, J. Behrmann, D. K. Duvenaud, and J.-H. Jacobsen, "Residual flows for invertible generative modeling," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 9916–9926.

[26] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 6572–6583.

[27] T. Chen, X. Zhai, M. Ritter, M. Lucic, and N. Houlsby, "Self-supervised GANs via auxiliary rotation loss," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 12146–12155.

[28] X. Chen *et al.*, "Variational lossy autoencoder," in *Proc. 5th Int. Conf. Learn. Representations*, 2017.

[29] X. Chen, N. Mishra, M. Rohaninejad, and P. Abbeel, "PixelSNAIL: An improved autoregressive generative model," in *Proc. 35th Int. Conf. Mach. Learn.*, 2017, pp. 864–872.

[30] R. Child, "Very deep VAEs generalize autoregressive models and can outperform them on images," in *Proc. Int. Conf. Learn. Representations*, 2021.

[31] R. Child, S. Gray, A. Radford, and I. Sutskever, "Generating long sequences with sparse transformers," 2019, *arXiv:1904.10509*.

[32] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1800–1807.

[33] K. M. Choromanski *et al.*, "Rethinking attention with performers," in *Proc. Int. Conf. Learn. Representations*, 2021.

[34] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," in *Proc. NIPS 2014 Workshop Deep Learn.*, 2014.

[35] R. Cornish, A. Caterini, G. Deligiannidis, and A. Doucet, "Relaxing bijectivity constraints with continuously indexed normalising flows," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 2133–2143.

[36] C. Cremer, X. Li, and D. Duvenaud, "Inference suboptimality in variational autoencoders," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1078–1086.

[37] B. Dai and D. Wipf, "Diagnosing and enhancing VAE models," in *Proc. Int. Conf. Learn. Representations*, 2018.

[38] N. D. Cao, I. Titov, and W. Aziz, "Block neural autoregressive flow," in *Proc. 35th Conf. Uncertainty Artif. Intell.*, 2019, pp. 1263–1273.

[39] E. L. Denton, S. Chintala, A. Szlam, and R. Fergus, "Deep generative image models using a Laplacian pyramid of adversarial networks," in *Proc. 28th Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 1486–1494.

[40] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol.*, 2019, pp. 4171–4186.

[41] L. Dinh, D. Krueger, and Y. Bengio, "NICE: Non-linear independent components estimation," in *Proc. 3rd Int. Conf. Learn. Representations Workshop*, 2015.

[42] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using Real NVP," in *Proc. 5th Int. Conf. Learn. Representations*, 2017.

[43] L. Dinh, J. Sohl-Dickstein, R. Pascanu, and H. Larochelle, "A RAD approach to deep mixture models," in *Proc. Int. Conf. Learn. Representations Workshop*, 2019.

[44] A. Dosovitskiy and T. Brox, "Generating images with perceptual similarity metrics based on deep networks," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 658–666.

[45] Y. Du and I. Mordatch, "Implicit generation and generalization in energy-based models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 3608–3618.

[46] L. L. Duan, "Transport monte carlo: High-accuracy posterior approximation via random transport," 2019, *arXiv:1907.10448*.

[47] E. Dupont, A. Doucet, and Y. W. Teh, "Augmented neural ODEs," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 3140–3150.

[48] E. Dupont, Y. W. Teh, and A. Doucet, "Generative models as distributions of functions," 2021, *arXiv:2102.04776*.

[49] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Cubic-spline flows," in *Proc. Int. Conf. Mach. Learn. Workshop*, 2019.

[50] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Neural spline flows," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 7511–7522.

[51] C. Durkan, I. Murray, and G. Papamakarios, "On contrastive learning for likelihood-free inference," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 2771–2781.

[52] P. Esser, R. Rombach, and B. Ommer, "Taming transformers for high-resolution image synthesis," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 12873–12883.

[53] C. Finlay, J.-H. Jacobsen, L. Nurbekyan, and A. Oberman, "How to train your neural ODE: The world of jacobian and kinetic regularization," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 3154–3164.

[54] R. Gao, E. Nijkamp, D. P. Kingma, Z. Xu, A. M. Dai, and Y. N. Wu, "Flow contrastive estimation of energy-based models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7515–7525.

[55] R. Gao, Y. Song, B. Poole, Y. N. Wu, and D. P. Kingma, "Learning energy-based models by diffusion recovery likelihood," in *Proc. Int. Conf. Learn. Representations*, 2021.

[56] M. Germain, K. Gregor, I. Murray, and H. Larochelle, "MADE: Masked autoencoder for distribution estimation," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 881–889.

[57] P. Ghosh, M. S. M. Sajjadi, A. Vergari, M. Black, and B. Scholkopf, "From variational to deterministic autoencoders," in *Proc. Int. Conf. Learn. Representations*, 2020.

[58] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. 27th Int. Conf. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.

[59] H. Gouk, E. Frank, B. Pfahringer, and M. J. Cree, "Regularisation of neural networks by enforcing Lipschitz continuity," *Mach. Learn.*, vol. 110, no. 2, pp. 393–416, 2021.

[60] K. Goyal, C. Dyer, and T. Berg-Kirkpatrick , "Exposing the implicit energy networks behind masked language models via metropolis–hastings," 2021, *arXiv:2106.02736*.

[61] W. Grathwohl, R. T. Q. Chen, J. Bettencourt, I. Sutskever, and D. Duvenaud, "FFJORD: Free-form continuous dynamics for scalable reversible generative models," in *Proc. Int. Conf. Learn. Representations*, 2019.

[62] W. Grathwohl, K. Swersky, M. Hashemi, D. Duvenaud, and C. J. Maddison, "Oops I took a gradient: Scalable sampling for discrete distributions," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 3831–3841.

[63] W. Grathwohl, K.-C. Wang, J.-H. Jacobsen, D. Duvenaud, and R. Zemel, "Learning the stein discrepancy for training and evaluating energy-based models without sampling," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 3732–3747.

[64] W. Grathwohl, K.-C. Wang, J.-H. Jacobsen, D. Duvenaud, M. Norouzi, and K. Swersky, "Your classifier is secretly an energy based model and you should treat it like one," in *Proc. Int. Conf. Learn. Representations*, 2020.

[65] W. S. Grathwohl, J. J. Kelly, M. Hashemi, M. Norouzi, K. Swersky, and D. Duvenaud, "No MCMC for me: Amortized sampling for fast and stable training of energy-based models," in *Proc. Int. Conf. Learn. Representations*, 2021.

[66] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra, "DRAW: A recurrent neural network for image generation," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 1462–1471.

[67] P. Grnarova, Y. Kilcher, K. Y. Levy, A. Lucchi, and T. Hofmann, "Generative minimization networks: Training GANs without competition," 2021, *arXiv:2103.12685*.

[68] P. Grnarova *et al.*, "A domain agnostic measure for monitoring and evaluating GANs," in *Prc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 12092–12102.

[69] A. Grover, M. Dhar, and S. Ermon, "Flow-GAN: Combining maximum likelihood and adversarial learning in generative models," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018, pp. 3069–3076.

[70] J. Gui, Z. Sun, Y. Wen, D. Tao, and J. Ye, "A review on generative adversarial networks: Algorithms, theory, and applications," 2020, *arXiv:2001.06937*.

[71] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C Courville, "Improved training of wasserstein GANs," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5769–5779.

[72] I. Gulrajani *et al.*, "PixelVAE: A latent variable model for natural images," in *Proc. Int. Conf. Learn. Representations*, 2017.

[73] I. Gulrajani, C. Raffel, and L. Metz, "Towards GAN benchmarks which require generalization," in *Proc. 7th Int. Conf. Learn. Representations*, 2019.

[74] M. Gutmann and A. Hyvärinen, "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2010, pp. 297–304.

[75] D. Ha and J. Schmidhuber, "World models," 2018, *arXiv:1803.10122*.

[76] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi, "Dream to control: Learning behaviors by latent imagination," in *Proc. Int. Conf. Learn. Representations*, 2020.

[77] T. Han, Y. Lu, S.-C. Zhu, and Y. N. Wu, "Alternating back-propagation for generator network," in *Proc. 31st AAAI Conf. Artif. Intell.*, 2017, pp. 1976–1984.

[78] L. Hasenclever, J. M. Tomczak, and M. Welling, "Variational inference with orthogonal normalizing flows," in *Proc. Workshop Bayesian Deep Learn. Neural Inf. Process. Syst.*, 2017.

[79] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.

[80] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local Nash equilibrium," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6629–6640.

[81] I. Higgins *et al.*, "beta-VAE: Learning basic visual concepts with a constrained variational framework," in *Proc. Int. Conf. Learn. Representations*, 2017.

[82] G. E. Hinton and T. J. Sejnowski, "Optimal perceptual inference," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 1983, pp. 448–453.

[83] G. E. Hinton, "Training products of experts by minimizing contrastive divergence," *Neural Comput.*, vol. 14, no. 8, pp. 1771–1800, 2002.

[84] G. E. Hinton, S. Osindero, and Y.-W. Teh, "A fast learning algorithm for deep belief nets," *Neural Comput.*, vol. 18, no. 7, pp. 1527–1554, 2006.

[85] J. Ho, X. Chen, A. Srinivas, Y. Duan, and P. Abbeel, "Flow++: Improving flow-based generative models with variational dequantization and architecture design," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 2722–2730.

[86] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Adv. Neural Inf. Process. Syst.* vol. 33, 2020, pp. 6840–6851.

[87] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.

[88] M. Hoffman, P. Sountsov, J. V. Dillon, I. Langmore, D. Tran, and S. Vasudevan, "NeuTra-lizing bad geometry in hamiltonian Monte Carlo using neural transport," 2019, *arXiv:1903.03704*.

[89] M. D. Hoffman and M. J. Johnson, "ELBO surgery: Yet another way to carve up the variational evidence lower bound," in *Proc. Conf. Neural Inf. Process. Syst. Workshop*, 2016.

[90] E. Hoogeboom, R. van den Berg, and M. Welling, "Emerging convolutions for generative normalizing flows," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 2771–2780.

[91] E. Hoogeboom, V. G. Satorras, J. Tomczak, and M. Welling, "The convolution exponential and generalized sylvester flows," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 18249–18260.

[92] E. Hoogeboom, D. Nielsen, P. Jaini, P. Forré, and M. Welling, "Argmax flows and multinomial diffusion: Learning categorical distributions," 2021, *arXiv:2102.05379*.

[93] E. Hoogeboom, J. Peters, R. van den Berg, and M. Welling, "Integer discrete flows and lossless compression," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 12134–12144.

[94] C.-W. Huang, L. Dinh, and A. Courville, "Augmented normalizing flows: Bridging the gap between generative flows and latent variable models," *Workshop Integr. Deep Neural Models Differ. Equ.*, 2020.

[95] C.-W. Huang, D. Krueger, A. Lacoste, and A. Courville, "Neural autoregressive flows," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 2078–2087.

[96] C.-W. Huang *et al.*, "Learnable explicit density for continuous latent space and variational inference," 2017, *arXiv:1710.02248*.

[97] H. Huang, Z. li, R. He, Z. Sun, and T. Tan, "IntroVAE: Introspective variational autoencoders for photographic image synthesis," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, 2018, pp. 52–63.

[98] X. Huang, Y. Li, O. Poursaeed, J. Hopcroft, and S. Belongie, "Stacked generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1866–1875.

[99] F. Huszár, "How (not) to train your generative model: Scheduled sampling, likelihood, adversary?," 2015, *arXiv:1511.05101*.

[100] A. Hyvärinen, "Estimation of non-normalized statistical models by score matching," *J. Mach. Learn. Res.*, vol. 6, pp. 695–709, 2005.

[101] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 448–456.

[102] P. Jaini, K. A. Selby, and Y. Yu, "Sum-of-squares polynomial flow," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 3009–3018.

[103] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with Gumbel-Softmax," in *Proc. Int. Conf. Learn. Representations*, 2017.

[104] Y. Jiang, S. Chang, and Z. Wang, "TransGAN: Two transformers can make one strong GAN," 2021, *arXiv:2102.07074*.

[105] L. Jing and Y. Tian, "Self-supervised visual feature learning with deep neural networks: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, May 04, 2020, doi: 10.1109/TPAMI.2020.2992393.

[106] A. Jolicoeur-Martineau, "The relativistic discriminator: A key element missing from standard GAN," in *Proc. Int. Conf. Learn. Representations*, 2018.

[107] A. Jolicoeur-Martineau, K. Li, R. Piché-Taillefer, T. Kachman, and I. Mitliagkas, "Gotta go fast when generating data with score-based models," 2021, *arXiv:2105.14080*.

[108] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, "Introduction to variational methods for graphical models," *Mach. Learn.*, vol. 37, no. 2, pp. 183–233, 1999.

[109] H. Jun *et al.*, "Distribution augmentation for generative modeling," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 5006–5019.

[110] L. Kaiser *et al.*, "Fast decoding in sequence models using discrete latent variables," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 2390–2399.

[111] M. Karami, D. Schuurmans, J. Sohl-Dickstein, L. Dinh, and D. Duckworth, "Invertible convolutional flow," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 5635–5645.

[112] A. Karnewar and O. Wang, "MSG-GAN: Multi-scale gradients for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 7796–7805.

[113] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *Proc. Int. Conf. Learn. Representations*, 2018.

[114] T. Karras, M. Aittala, J. Hellsten, S. Laine, J. Lehtinen, and T. Aila, "Training generative adversarial networks with limited data," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 12104–12114.

[115] T. Karras, M. Aittala, S. Laine, E. Härkönen, J. Hellsten, J. Lehtinen, and T. Aila, "Alias-free generative adversarial networks," 2021, *arXiv:2106.12423*.

[116] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4396–4405.

[117] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of StyleGAN," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8107–8116.

[118] A. Katharopoulos, A. Vyas, N. Pappas, and F. Fleuret, "Transformers are RNNs: Fast autoregressive transformers with linear attention," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 5156–5165.

[119] H. Kim, G. Papamakarios, and A. Mnih, "The Lipschitz constant of self-attention," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 5562–5571.

[120] T. Kim and Y. Bengio, "Deep directed generative models with energy-based probability estimation," 2016, *arXiv:1606.03439*.

[121] Y. Kim, S. Wiseman, A. Miller, D. Sontag, and A. Rush, "Semi-amortized variational autoencoders," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 2678–2687.

[122] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. Int. Conf. Learn. Representations*, 2014.

[123] D. P. Kingma and P. Dhariwal, "Glow: Generative flow with invertible 1x1 convolutions," in *Proc. Conf. Neural Inf. Process. Syst.*, 2018, pp. 10215–10224.

[124] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, "Improved variational inference with inverse autoregressive flow," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 4743–4751.

[125] I. Kobyzev, S. Prince, and M. Brubaker, "Normalizing flows: An introduction and review of current methods," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, May 07, 2020, doi: 10.1109/TPAMI.2020.2992934.

[126] A. Krizhevsky, "Learning multiple layers of features from tiny images," Univ. Toronto, Toronto, ON, Canada, Tech. Rep. Krizhevsky_2009_17719, 2009.

[127] K. Kumar *et al.*, "MelGAN: Generative adversarial networks for conditional waveform synthesis," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 14910–14921.

[128] M. Kumar *et al.*, "VideoFlow: A flow-based generative model for video," in *Proc. Int. Conf. Mach. Learn. Workshop*, 2019.

[129] R. Kumar, S. Ozair, A. Goyal, A. Courville, and Y. Bengio, "Maximum entropy generators for energy-based models," 2019, *arXiv:1901.08508*.

[130] H. Larochelle and I. Murray, "The neural autoregressive distribution estimator," in *Proc. 14th Int. Conf. Artif. Intell. Statist.*, 2011, pp. 29–37.

[131] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther, "Autoencoding beyond pixels using a learned similarity metric," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 1558–1566.

[132] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. J. Huang, "A tutorial on energy-based learning," *Predicting Structured Data 1*, 2006.

[133] A. X. Lee, R. Zhang, F. Ebert, P. Abbeel, C. Finn, and S. Levine, "Stochastic adversarial video prediction," 2018, *arXiv:1804.01523*.

[134] Z. Li, Y. Chen, and F. T. Sommer, "A neural network MCMC sampler that maximizes proposal entropy," *Entropy*, vol. 23, no. 3, p. 269, 2021.

[135] J. H. Lim and J. C. Ye, "Geometric GAN," 2017, *arXiv:1705.02894*.

[136] J. Lin, R. Zhang, F. Ganz, S. Han, and J.-Y. Zhu, "Anycost GANs for interactive image synthesis and editing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 14986–14996.

[137] B. Liu, Y. Zhu, K. Song, and A. Elgammal, "Towards faster and stabilized GAN training for high-fidelity few-shot image synthesis," in *Proc. Int. Conf. Learn. Representations*, 2021.

[138] K. Liu, G. Qiu, W. Tang, and F. Zhou, "Spectral regularization for combating mode collapse in GANs," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 6381–6389.
[139] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised image-to-image translation networks," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 700–708.
[140] X. Liu and C.-J. Hsieh, "Rob-GAN: Generator, discriminator, and adversarial attacker," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 11226–11235.
[141] G. Loaiza-Ganem and J. P. Cunningham, "The continuous Bernoulli: Fixing a pervasive error in variational autoencoders," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 13266–13276.
[142] C. Lu, J. Chen, C. Li, Q. Wang, and J. Zhu, "Implicit normalizing flows," in *Proc. Int. Conf. Learn. Representations*, 2021, pp. 13266–13276.
[143] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet, "Are GANs created equal? A large-scale study," in *Proc. Conf. Neural Inf. Process. Syst.*, 2018, pp. 698–707.
[144] X. Ma, X. Kong, S. Zhang, and E. Hovy, "MaCow: Masked convolutional generative flow," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 5893–5902.
[145] X. Ma, C. Zhou, and E. Hovy, "MAE: Mutual posterior-divergence regularization for variational AutoEncoders," in *Proc. Int. Conf. Learn. Representations*, 2019.
[146] L. Maaløe, M. Fraccaro, V. Liévin, and O. Winther, "BIVA: A very deep hierarchy of latent variables for generative modeling," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 6551–6562.
[147] C. J. Maddison, A. Mnih, and Y. W. Teh, "The concrete distribution: A continuous relaxation of discrete random variables," in *Proc. Int. Conf. Learn. Representations*, 2017.
[148] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *Proc. Int. Conf. Learn. Representations*, 2018.
[149] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, "Adversarial autoencoders," 2016, *arXiv:1511.05644*.
[150] X. Mao, Q. Li, H. Xie, R. Y. K. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2813–2821.
[151] S. Mehri *et al.*, "SampleRNN: An unconditional end-to-end neural audio generation model," in *Proc. Int. Conf. Learn. Representations*, 2017.
[152] C. Meng, J. Song, Y. Song, S. Zhao, and S. Ermon, "Improved autoregressive modeling with distribution smoothing," in *Proc. Int. Conf. Learn. Representations*, 2021.
[153] C. Meng, L. Yu, Y. Song, J. Song, and S. Ermon, "Autoregressive score matching," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 6673–6683.
[154] J. Menick and N. Kalchbrenner, "Generating high fidelity images with subscale pixel networks and multidimensional upscaling," in *Proc. Int. Conf. Learn. Representations*, 2019.
[155] M. Mirza and S. Osindero, "Conditional generative adversarial nets," 2014, *arXiv:1411.1784*.
[156] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral normalization for generative adversarial networks," in *Proc. Int. Conf. Learn. Representations*, 2018.
[157] T. Müller, B. Mcwilliams, F. Rousselle, M. Gross, and J. Novák, "Neural importance sampling," *ACM Trans. Graph.*, vol. 38, no. 5, pp. 145:1–145:19, 2019.
[158] C. Nash and C. Durkan, "Autoregressive energy machines," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 1735–1744.
[159] K. Neklyudov, E. Egorov, and D. P. Vetrov, "The implicit Metropolis-Hastings algorithm," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 13955–13965.
[160] M. Ngxande, J. Tapamo, and M. Burke, "DepthwiseGANs: Fast training generative adversarial networks for realistic image synthesis," in *Proc. Southern Afr. Univ. Power Eng. Conf./Robot. Mechatronics/Pattern Recognit. Assoc. South Afr.*, 2019, pp. 111–116.
[161] A. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," 2021, *arXiv:2102.09672*.
[162] W. Nie, N. Narodytska, and A. Patel, "RelGAN: Relational generative adversarial networks for text generation," in *Proc. Int. Conf. Learn. Representations*, 2019.
[163] D. Nielsen, P. Jaini, E. Hoogeboom, O. Winther, and M. Welling, "SurVAE flows: Surjections to bridge the gap between VAEs and flows," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 12685–12696.
[164] E. Nijkamp *et al.*, "Learning energy-based model with flow-based backbone by neural transport MCMC," 2020, *arXiv:2006.06897*.
[165] E. Nijkamp, M. Hill, T. Han, S.-C. Zhu, and Y. N. Wu, "On the anatomy of MCMC-based maximum likelihood learning of energy-based models," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, no. 4, 2020, pp. 5272–5280.
[166] E. Nijkamp, M. Hill, S.-C. Zhu, and Y. N. Wu, "Learning non-convergent non-persistent short-run MCMC toward energy-based model," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 5232–5242.
[167] E. Nijkamp, B. Pang, T. Han, L. Zhou, S.-C. Zhu, and Y. N. Wu, "Learning multi-layer latent variable model via variational optimization of short run MCMC for approximate inference," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 361–378.
[168] S. Norouzi, D. J. Fleet, and M. Norouzi, "Exemplar VAE: Linking generative models, nearest neighbor retrieval, and data augmentation," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 8753–8764.
[169] S. Nowozin, B. Cseke, and R. Tomioka, "f-GAN: Training generative neural samplers using variational divergence minimization," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 271–279.
[170] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier GANs," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 2642–2651.
[171] D. Onken, S. W. Fung, X. Li, and L. Ruthotto, "OT-Flow: Fast and accurate continuous normalizing flows via optimal transport," in *Proc. AAAI Conf. Artif. Intell.*, 2021, pp. 9223–9232.
[172] G. Ostrovski, W. Dabney, and R. Munos, "Autoregressive Quantile networks for generative modeling," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 3936–3945.
[173] A. Oussidi and A. Elhassouny, "Deep generative models: Survey," in *Proc. Int. Conf. Intell. Syst. Comput. Vis.*, 2018, pp. 1–8.
[174] T. Pang, K. Xu, C. Li, Y. Song, S. Ermon, and J. Zhu, "Efficient learning of generative models via finite-difference score matching," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, vol. 34, 2020.
[175] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, "Normalizing flows for probabilistic modeling and inference," *J. Mach. Learn. Res.*, vol. 22, no. 57, pp. 1–64, 2021.
[176] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked autoregressive flow for density estimation," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 2335–2344.
[177] N. Parmar *et al.*, "Image transformer," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 4055–4064.
[178] L. S. Pontryagin, *Mathematical Theory of Optimal Processes*. Evanston, IL, USA: Routledge, 2018.
[179] R. Prenger, R. Valle, and B. Catanzaro, "WaveGlow: A flow-based generative network for speech synthesis," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2019, pp. 3617–3621.
[180] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *Proc. Int. Conf. Learn. Representations*, 2016.
[181] A. Ramesh *et al.*, "Zero-shot text-to-image generation," 2021, in *Proc. 38th Int. Conf. Mach. Learn.*, 2021, pp. 8821–8831.
[182] A. Razavi, A. van den Oord, and O. Vinyals, "Generating diverse high-fidelity images with VQ-VAE-2," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 14837–14847.
[183] S. Reed *et al.*, "Parallel multiscale autoregressive density estimation," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 2912–2921.
[184] D. J. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 1530–1538.
[185] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014, pp. 1278–1286.
[186] G. O. Roberts and R. L. Tweedie, "Exponential convergence of Langevin distributions and their discrete approximations," *Bernoulli*, vol. 2, no. 4, pp. 341–363, 1996.
[187] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training GANs," in *Proc. Conf. Neural Inf. Process. Syst.*, 2016, pp. 2234–2242.
[188] T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma, "PixelCNN++: Improving the PixelCNN with discretized logistic mixture likelihood and other modifications," in *Proc. Int. Conf. Learn. Representations*, 2017.
[189] T. Salimans, D. Kingma, and M. Welling, "Markov chain Monte Carlo and variational inference: Bridging the gap," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 1218–1226.

[190] S. Saremi, A. Mehrjou, B. Schölkopf, and A. Hyvärinen, "Deep energy estimator networks," 2018, *arXiv:1805.08306*.

[191] H. Sasaki, C. G. Willcocks, and T. P. Breckon, "UNIT-DDPM: UNpaired image translation with denoising diffusion probabilistic models," 2021, *arXiv:2104.05358*.

[192] J. Schmidhuber, "Making the world differentiable: On using self-supervised fully recurrent neural networks for dynamic reinforcement learning and planning in non-stationary environments," Technische Universitaet Muenchen, Munich, Bavaria, Germany, Tech. Rep. FKI-126–90, 1990.

[193] J. Schmidhuber, "Generative adversarial networks are special cases of artificial curiosity (1990) and also closely related to predictability minimization (1991)," *Neural Netw.*, vol. 127, pp. 58–66, 2020.

[194] J. Shi, S. Sun, and J. Zhu, "A spectral approach to gradient estimation for implicit distributions," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 4644–4653.

[195] S. Sinha, H. Zhang, A. Goyal, Y. Bengio, H. Larochelle, and A. Odena, "Small-GAN: Speeding up GAN training using coresets," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 9005–9015.

[196] V. Sitzmann, J. N. P. Martel, A. W. Bergman, D. B. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," Advances in *Proc. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 7462–7473.

[197] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep Unsupervised Learning using Nonequilibrium Thermodynamics," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 2256–2265.

[198] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *Proc. Int. Conf. Learn. Representations*, 2021.

[199] J. Song, S. Zhao, and S. Ermon, "A-NICE-MC: Adversarial training for MCMC," in *Proc. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5140–5150.

[200] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 11895–11907.

[201] Y. Song, S. Garg, J. Shi, and S. Ermon, "Sliced score matching: A scalable approach to density and score estimation," in *Proc. Conf. Uncertainty Artif. Intell.*, 2020, pp. 574–584.

[202] Y. Song and D. P. Kingma, "How to train your energy-based models," 2021, *arXiv:2101.03288*.

[203] Y. Song, C. Meng, and S. Ermon, "MintNet: Building invertible neural networks with masked convolutions," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 11004–11014.

[204] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *Proc. Int. Conf. Learn. Representations*, 2021.

[205] Y. Song, Q. Ye, M. Xu, and T.-Y. Liu, "Discriminator contrastive divergence: Semi-amortized generative modeling by exploring energy of the discriminator," 2020, *arXiv:2004.01704*.

[206] I. Sutskever and T. Tieleman, "On the convergence properties of contrastive divergence," in *Proc. 13th Int. Conf. Artif. Intell. Statist.*, 2010, pp. 789–795.

[207] K. Swersky, M. Ranzato, D. Buchman, B. M. Marlin, and N. de Freitas, "On autoencoders and score matching for energy based models," in *Proc. 28th Int. Conf. Mach. Learn.*, 2011, pp. 1201–1208.

[208] C. K. Sønderby, J. Caballero, L. Theis, W. Shi, and F. Huszár, "Amortised MAP inference for image super-resolution," in *Proc. Int. Conf. Learn. Representations*, 2017.

[209] C. K. Sønderby, T. Raiko, L. Maaløe, S. K. Sønderby, and O. Winther, "Ladder variational autoencoders," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 3745–3753.

[210] M. Tancik *et al.*, "Fourier features let networks learn high frequency functions in low dimensional domains," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 7537–7547.

[211] H. Thanh-Tung and T. Tran, "Catastrophic forgetting and mode collapse in GANs," in *Proc. Int. Joint Conf. Neural Netw.*, 2020, pp. 1–10.

[212] L. Theis and M. Bethge, "Generative image modeling using spatial LSTMs," in *Proc. 28th Conf. Neural Inf. Process. Syst.*, 2015, pp. 1927–1935.

[213] L. Theis, A. van den Oord, and M. Bethge, "A note on the evaluation of generative models," 2016, in *Proc. Int. Conf. Learn. Representations*, 2016.

[214] T. Tieleman, "Training restricted Boltzmann machines using approximations to the likelihood gradient," in *Proc. 25th Int. Conf. Mach. Learn.*, 2008, pp. 1064–1071.

[215] M. Titsias and P. Dellaportas, "Gradient-based adaptive Markov chain Monte Carlo," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 15730–15739.

[216] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf, "Wasserstein auto-encoders," in *Int. Conf. Learn. Representations*, 2018.

[217] J. Tomczak and M. Welling, "VAE with a VampPrior," in *Proc. 21st Int. Conf. Artif. Intell. Statist.*, 2018, pp. 1214–1223.

[218] D. Tran, R. Ranganath, and D. M. Blei, "Variational Gaussian process," in *Proc. Int. Conf. Learn. Representations*, 2016.

[219] D. Tran, K. Vafa, K. Agrawal, L. Dinh, and B. Poole, "Discrete flows: Invertible generative models of discrete data," in *Proc. 33rd Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 14719–14728.

[220] N.-T. Tran, V.-H. Tran, N.-B. Nguyen, T.-K. Nguyen, and N.-M. Cheung, "On data augmentation for GAN training," *IEEE Trans. Image Process.*, vol. 30, pp. 1882–1897, Jan. 2021.

[221] N.-T. Tran, V.-H. Tran, B.-N. Nguyen, L. Yang, and N.-M. M. Cheung, "Self-supervised GAN: Analysis and improvement with multi-class minimax game," in *Proc. Conf. Neural Inf. Process. Syst.*, 2019, pp. 13253–13264.

[222] R. E. Turner and M. Sahani, "Two problems with variational expectation maximisation for time series models," in *Bayesian Time Series Models*. Cambridge, U.K.: Cambridge Univ. Press, 2011, pp. 104–124.

[223] R. Turner, J. Hung, E. Frank, Y. Saatchi, and J. Yosinski, "Metropolis-Hastings generative adversarial networks," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6345–6353.

[224] B. Tzen and M. Raginsky, "Theoretical guarantees for sampling and inference in generative models with latent diffusions," in *Proc. Annu. Conf. Learn. Theory*, 2019, pp. 3084–3114.

[225] B. Uria, I. Murray, and H. Larochelle, "RNADE: The real-valued neural autoregressive density-estimator," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.*, 2013, pp. 2175–2183.

[226] A. Vahdat and J. Kautz, "NVAE: A deep hierarchical variational autoencoder," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 19667–19679.

[227] A. Vahdat, K. Kreis, and J. Kautz, "Score-based generative modeling in latent space," 2021, *arXiv:2106.05931*.

[228] G. Valvano, A. Leo, and S. A. Tsaftaris, "Learning to segment from scribbles using multi-scale adversarial attention gates," *IEEE Trans. Med. Imag.*, vol. 40, no. 8, pp. 1990–2001, Aug. 2021.

[229] A. van den Oord *et al.*, "WaveNet: A generative model for raw audio," 2016, *arXiv:1609.03499*.

[230] A. van den Oord *et al.*, "Parallel WaveNet: Fast high-fidelity speech synthesis," in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 3918–3926.

[231] A. van den Oord, N. Kalchbrenner, L. Espeholt, K. kavukcuoglu, O. Vinyals, and A. Graves, "Conditional image generation with PixelCNN decoders," in *Proc. Conf. Neural Inf. Process. Syst.*, 2016, pp. 4797–4805.

[232] A. Van Den Oord, N. Kalchbrenner, and K. Kavukcuoglu, "Pixel recurrent neural networks," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 1747–1756.

[233] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6309–6318.

[234] A. Vaswani *et al.*, "Attention is all you need," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.

[235] C. Villani, *Topics in Optimal Transportation*. Providence, RI, USA: Amer. Math. Soc., 2003.

[236] C. Villani, *Optimal Transport: Old and New*. Berlin, Germany: Springer, 2008.

[237] P. Vincent, "A connection between score matching and denoising autoencoders," *Neural Comput.*, vol. 23, no. 7, pp. 1661–1674, 2011.

[238] A. Wang and K. Cho, "BERT has a mouth, and it must speak: BERT as a Markov random field language model," in *Proc. Workshop Methods Optimizing Evaluating Neural Lang. Gener.*, 2019, pp. 30–36.

[239] S. Wang, B. Z. Li, M. Khabsa, H. Fang, and H. Ma, "Linformer: Self-attention with linear complexity," 2020, *arXiv:2006.04768*.

[240] Y. Wang, G. Zhang, and J. Ba, "On solving minimax optimization locally: A follow-the-ridge approach," in *Proc. Int. Conf. Learn. Representations*, 2020.

[241] D. Watson, J. Ho, M. Norouzi, and W. Chan, "Learning to efficiently sample from diffusion probabilistic models," 2021, *arXiv:2106.03802*.

[242] A. Wehenkel and G. Louppe, "Unconstrained Monotonic neural networks," in *Proc. 33rd Conf. Neural Inf. Process. Syst.*, 2019, pp. 1545–1555.
[243] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *Proc. 28th Int. Conf. Mach. Learn.*, 2011, pp. 681–688.
[244] H. Wu, J. Köhler, and F. Noé, "Stochastic normalizing flows," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 5933–5944.
[245] Z. Xiao, K. Kreis, J. Kautz, and A. Vahdat, "VAEBM: A symbiosis between variational autoencoders and energy-based models," in *Proc. Int. Conf. Learn. Representations*, 2021.
[246] J. Xie, Y. Lu, R. Gao, S. Zhu, and Y. N. Wu, "Cooperative training of descriptor and generator networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 1, pp. 27–45, Jan. 2020.
[247] J. Xie, Y. Lu, S.-C. Zhu, and Y. N. Wu, "A theory of generative ConvNet," in *Proc. 33rd Int. Conf. Mach. Learn.*, 2016, pp. 2635–2644.
[248] L. Yang and G. E. Karniadakis, "Potential flow generator with L_2 optimal transport regularity for generative models," *IEEE Trans. Neural Netw. Learn. Syst.*, early access, Oct. 20, 2020, doi: 10.1109/TNNLS.2020.3028042.
[249] X. Yi, E. Walia, and P. Babyn, "Generative adversarial network in medical imaging: A review," *Med. Image Anal.*, vol. 58, 2019, Art. no. 101552.
[250] L. Yu, Y. Song, J. Song, and S. Ermon, "Training deep energy-based models with f-divergence minimization," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 10957–10967.
[251] S. Zagoruyko and N. Komodakis, "Wide residual networks," in *Proc. British Mach. Vis. Conf.*, 2016.
[252] H. Zhang, X. Gao, J. Unterman, and T. Arodz, "Approximation capabilities of neural ordinary differential equations," 2019, *arXiv:1907.12998*.
[253] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, "Self-attention generative adversarial networks," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 7354–7363.
[254] H. Zhang *et al.*, "StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5908–5916.
[255] L. Zhang, W. E., and L. Wang, "Monge-Ampère flow for generative modeling," 2018, *arXiv:1809.10188*.
[256] J. Zhao, M. Mathieu, and Y. LeCun , "Energy-based generative adversarial network," in *Proc. Int. Conf. Learn. Representations*, 2017.
[257] S. Zhao, J. Song, and S. Ermon, "Learning hierarchical features from deep generative models," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 4091–4099.
[258] S. Zhao, J. Song, and S. Ermon, "Towards deeper understanding of variational autoencoding models," 2017, pp. *arXiv:1702.08658*.
[259] S. Zhao, J. Song, and S. Ermon, "InfoVAE: Balancing learning and inference in variational autoencoders," in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 5885–5892.
[260] S. Zhao, Z. Liu, J. Lin, J.-Y. Zhu, and S. Han, "Differentiable augmentation for data-efficient GAN training," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020, pp. 7559–7570.
[261] Z. Zhao, Z. Zhang, T. Chen, S. Singh, and H. Zhang, "Image augmentations for GAN training," 2020, *arXiv:2006.02595*.
[262] J. Zhong, X. Liu, and C.-J. Hsieh, "Improving the speed and quality of GAN by adversarial training," 2020, *arXiv:2008.03364*.
[263] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2242–2251.
[264] Z. Ziegler and A. Rush, "Latent normalizing flows for discrete sequences," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 7673–7682.
[265] B. Pang, T. Han, E. Nijkamp, S.-C. Zhu, and Y. N. Wu, "Learning latent space energy-based prior model," in *Proc. Conf. Neural Inf. Process. Syst.*, 2020.

Sam Bond-Taylor is currently working toward the PhD degree at the Department of Computer Science, Durham University, Durham, U.K. His research interests include unsupervised deep learning methods and the connections with human learning. In particular, the development of generative models and machine reasoning systems. He is a teaching assistant on the deep learning module at Durham University.

Adam Leach is currently working toward the PhD degree at the Department of Computer Science, Durham University, Durham, U.K. His research focuses on applying deep generative models and reinforcement learning techniques to molecular modelling problems such as protein folding and docking. He is a teaching assistant on the reinforcement learning module at Durham University.

Yang Long is currently an assistant professor with the Department of Computer Science, Durham University. He is also an MRC innovation fellow aiming to design scalable AI solutions for large-scale healthcare applications. His research background is in the highly interdisciplinary field of computer vision and machine learning. He has authored/co-authored more than 20 top-tier papers in refereed journals/conferences such as the *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *IEEE Transactions on Image Processing*, CVPR, AAAI, and ACM MM.

Chris G. Willcocks is currently an assistant professor in computer science with Durham University, where his interdisciplinary research focuses on generative models, medical image computing, computational biophysics, and machine reasoning. He teaches deep learning, reinforcement learning, and cyber security, and publishes top-tier journal and conference papers in venues such as ICLR, *Physical Review X*, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *IEEE Transactions on Medical Imaging*, and *IEEE Transactions on Information Forensics and Security*.