
Funnels

Exact maximum likelihood with dimensionality reduction

Samuel Klein
University of Geneva
samuel.klein@unige.ch

John A. Raine
University of Geneva
john.raine@unige.ch

Sebastian Pina-Otey
University of Geneva
sebastian.pinaotey@unige.ch

Slava Voloshynovskiy
University of Geneva
svolos@unige.ch

Tobias Golling
University of Geneva
tobias.golling@unige.ch

Abstract

Normalizing flows are diffeomorphic, typically dimension-preserving, models trained using the likelihood of the model. We use the SurVAE framework to construct dimension reducing surjective flows via a new layer, known as the funnel. We demonstrate its efficacy on a variety of datasets, and show it improves upon or matches the performance of existing flows while having a reduced latent space size. The funnel layer can be constructed from a wide range of transformations including restricted convolution and feed forward layers.

1 Introduction

Normalizing flows [1–6] are powerful generative models that can be used to approximate the likelihood of data directly, unlike variational methods which calculate a lower bound [7, 8]. Flows are restricted to being diffeomorphisms and, unless the data is known to lie on a Riemannian manifold [6], preserve the dimensionality of the data. This leads to undesirable effects such as out of distribution samples being assigned higher likelihoods than inliers [9] and difficulties in modelling complex high dimensional data [1, 10]. These issues stem, at least in part, from the curse of dimensionality. Furthermore, autoregressive flows are the most flexible family of flows, and the cost of inverting these models scales linearly with the dimension of the data [6]. Therefore, methods for generalizing flows to be dimensionality reducing are required.

In contrast to strictly invertible functions, like those used in flows, surjective transformations can be used to exactly calculate the likelihood while changing the dimension of the data. The SurVAE framework can be used to build modular surjective transformations that allow for exact likelihood inference [11]. This framework provides a bridge between variational autoencoders (VAEs), which can perform arbitrary dimension reduction but only provide a lower bound for the likelihood, and flows which are typically dimension preserving diffeomorphisms. Within the SurVAE framework the log likelihood of data $\mathbf{x} \in \mathcal{X}$ is modelled as

$$\log(p(\mathbf{x})) = \log(p_\theta(\mathbf{z})) + \mathcal{V}(\mathbf{x}, \mathbf{z}) + \mathcal{E}(\mathbf{x}, \mathbf{z}), \quad \mathbf{z} \sim q(\mathbf{z}|\mathbf{x}), \quad \mathbf{z} \in \mathcal{Z},$$

where $q(\mathbf{z}|\mathbf{x})$ is an amortized variational distribution and p_θ is the base distribution defined on the latent space \mathcal{Z} . The \mathcal{V} term is referred to as the likelihood contribution and \mathcal{E} as the bound looseness.

The likelihood contribution for a deterministic surjection $f : \mathcal{X} \rightarrow \mathcal{Z}$ is given by

$$\mathcal{V}(\mathbf{x}, \mathbf{z}) = \lim_{q(\mathbf{z}|\mathbf{x}) \rightarrow \delta(\mathbf{z}-f(\mathbf{x}))} \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \frac{p(\mathbf{x}|\mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \right], \quad (1)$$

where $p(\mathbf{x}|\mathbf{z})$ is a stochastic generative process. The bound looseness term is zero if the function $f(\mathbf{x})$ satisfies the *right inverse condition* (A.1).

The SurVAE paper [11] provides useful novel layers but no principled method for directly extending flows. In this work we derive a layer which extends flows to (i) be surjective with exact likelihoods; (ii) have latent spaces with smaller dimensions than the data space and (iii) propagate compressed information explicitly. Our method also allows new types of invertible transformations to be derived, which we demonstrate by converting convolutional neural networks and multilayer perceptrons into efficient exact maximum likelihood estimators (A.7). Deriving such transformations allows flows to be extended to large dimensional datasets. We benchmark the proposed method on a wide variety of real world datasets and demonstrate their applicability and usefulness for downstream tasks. The code required to repeat the experiments is provided at https://github.com/sambklein/funnels_repo.

2 Exact maximum likelihood with dimension reduction

Given $\mathbf{x} \in \mathcal{X}$, any diffeomorphism F^{-1} that acts on a subset of the features in \mathbf{x} as

$$\mathbf{z} = F^{-1}(\mathbf{x}_{m:n}), \quad (2)$$

for $m < n = \dim(\mathcal{X})$, satisfies the *right inverse condition*. By taking F to be parameterized by a function $\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})$ a flexible new family of models can be derived. This transformation is a generalization of the tensor slicing layer [3, 11–13], where one of the key differences is F_θ is any function that can be inverted for $\mathbf{x}_{m:n}$ given $\mathbf{x}_{1:m-1}$ (A.3). It also explicitly allows the dependence of the model on $\mathbf{x}_{1:m-1}$ to be controlled and propagated efficiently.

The generative transformation within the SurVAE framework is given by

$$G(\mathbf{z}) = \begin{bmatrix} \mathbf{x}_{1:m-1} \sim p(\mathbf{x}_{1:m-1}|\mathbf{z}) \\ F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{z})}(\mathbf{z}|\mathbf{x}_{1:m-1}) \end{bmatrix}. \quad (3)$$

The process $p(\mathbf{x}_{1:m-1}|\mathbf{z})$ can be any continuous density that maps to the full range of $\mathbf{x}_{1:m-1}$, this can be parameterized by a density $p_\phi(\mathbf{x}_{1:m-1}|\mathbf{z})$. To derive the likelihood contribution for a funnel layer we start by following the SurVAE framework [11] and using the composition of a smooth function g with a Dirac δ -function and a diffeomorphism f where

$$\int d\mathbf{x} \delta(g(\mathbf{x})) f(g(\mathbf{x})) \left| \det \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}} \right| = \int d\mathbf{u} \delta(\mathbf{u}) f(\mathbf{u}), \quad (4)$$

to conclude that

$$\delta(g(\mathbf{x})) = \left| \det \frac{\partial g(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}_0}^{-1} \delta(\mathbf{x} - \mathbf{x}_0), \quad (5)$$

where \mathbf{x}_0 is the root of $g(\mathbf{x})$. This assumes the root is unique, the Jacobian is non-singular and f has compact support. Using this result we take $g(\mathbf{x}) = \mathbf{z} - F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})$ where the unique root is given by $\mathbf{x}_0 = F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}(\mathbf{z})$ and

$$\begin{aligned} q(\mathbf{z}|\mathbf{x}) &= \delta(\mathbf{z} - F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})) \\ &= \delta(\mathbf{x}_{m:n} - F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}(\mathbf{z})) |\det(J_\theta(\mathbf{x}_{m:n}))|^{-1}, \end{aligned}$$

where

$$J_\theta = \frac{\partial F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})}{\partial \mathbf{x}_{m:n}} \bigg|_{\mathbf{x}_{m:n}=F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}(\mathbf{z})}.$$

With a change of variables to $\mathbf{x}'_{m:n} = F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}(\mathbf{z})$ such that

$$\begin{aligned} d\mathbf{x}'_{m:n} &= d\mathbf{z} |\det(J_\theta^{-1}(\mathbf{x}_{m:n}))| \\ &= d\mathbf{z} |\det(J_\theta(\mathbf{x}_{m:n}))|^{-1}, \end{aligned}$$




Figure 1: Funnel layer with the forward (left) and inverse (right) passes. In each pass $\mathbf{x}_- = \mathbf{x}_{1:m-1}$ conditions the diffeomorphism F . In the inverse pass \mathbf{x}_- is first sampled from $p(\mathbf{x}_-|z)$ and then used to condition the inverse of the diffeomorphism to return a sample of $\mathbf{x}_+ = \mathbf{x}_{m:n}$.

Eq. (1) can be written as

$$\begin{aligned}
\mathcal{V}(\mathbf{x}) &= \int d\mathbf{z} \delta\left(\mathbf{z} - F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})\right) \log\left(\frac{p_\phi(\mathbf{x}_{1:m-1}|\mathbf{z})}{\delta\left(\mathbf{z} - F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})\right)}\right) \\
&= \int d\mathbf{x}'_{m:n} \delta(\mathbf{x}_{m:n} - \mathbf{x}'_{m:n}) \log(p_\phi(\mathbf{x}_{1:m-1}|\mathbf{z}) |\det(J_\theta(\mathbf{x}'_{m:n}))|) \\
&= \log\left(p_\phi\left(\mathbf{x}_{1:m-1} | F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})\right)\right) + \log(|\det(J_\theta(\mathbf{x}_{m:n}))|). \tag{6}
\end{aligned}$$

This expression can be understood as the determinant of the Jacobian from the subset of the coordinates for which the model is a diffeomorphism ($\mathbf{x}_{m:n}$), plus the log likelihood of the coordinates which condition the diffeomorphism ($\mathbf{x}_{1:m-1}$). The log likelihood acts as a reconstruction term that penalises moving away from invertibility. This is the derivation of the likelihood contribution for a funnel layer with an inference surjection; the same for a generative surjection can also be derived as shown in A.2.

The likelihood contribution in Eq. (6) can be used to form models that reduce dimensions by arbitrary amounts, can be trained on exact maximum likelihood estimation and permit exact latent variable inference. This likelihood contribution defines a *funnel* layer. The transformation F_θ used in this layer needs to be invertible for $\mathbf{x}_{m:n}$ given $\{\mathbf{x}_{1:m-1}, \mathbf{z}\}$ and have a tractable Jacobian determinant. Such transformations are abundantly available in the form of (conditional) invertible neural networks (INNs), but also appear in linear transformations such as convolutional kernels. A sketch of the layer can be seen in Fig. 1. The likelihood contribution in Eq. (6) can be used to define funnel layers.

3 Novel funnel layers

In this section we derive some novel funnel layers to demonstrate the use of Eq. (6) for deriving layers from existing methods.

Convolutional kernels. Convolutional kernels on one dimensional data can be turned into funnels as follows. Given a one dimensional vector $\mathbf{x} \in \mathbb{R}^n$ the forward transformation of the funnel can be factorised to form a 2×1 convolutional kernel $[a, b]$ with stride two

$$q(\mathbf{z}|\mathbf{x}) = \prod_{i=1}^{n/2} q(z_i | x_{2i-1}, x_{2i})$$

with

$$\begin{aligned}
z_i &= F_{\theta(x_{2i-1}, x_{2i})}^{-1}(x_{2i-1}) \\
&= ax_{2i-1} + bx_{2i}
\end{aligned}$$

where $[a, b]$ are the parameters of the convolutional kernel. Replacing $\mathbf{x}_{1:m-1}$ with x_{2i-1} in Eq. (6) the Jacobian of each kernel is given by

$$J_i = a,$$




Figure 2: A funnel NSF. The black line in the left figure is a true standard normal distribution and in blue is a histogram of the model encoding. Samples drawn from the prior (middle) and model likelihood (right).

and the total Jacobian is given by $\frac{n}{2}a$. The total likelihood contribution in this case is given by

$$\mathcal{V}(\mathbf{x}) = \log \left(p_\phi \left(\{x_{2i}\}_{i=1}^{n/2} | \mathbf{z} \right) \right) + \log \left(\frac{n}{2} |a| \right). \quad (7)$$

Given $x'_{2i} \sim p_\phi \left(\{x_{2i}\}_{i=1}^{n/2} | \mathbf{z} \right)$, the kernel $[a, b]$, and \mathbf{z} we can find

$$x_{2i-1} = J_i^{-1}(z_i - bx'_{2i}), \quad (8)$$

for all i such that the right inverse condition is satisfied. Here the different elements of the kernel play different roles in the layer, with a appearing directly in the loss and b implicitly in p_ϕ . Choosing x_{2i-1} to replace $\mathbf{x}_{m:n}$ was arbitrary. These considerations mean some features in the input vector are treated differently, partially breaking the spatial invariance of convolutional kernels. Details of how this is generalised can be found in A.7.1

MLP funnels. Non-invertible activation functions such as ReLU [14] can be used with exact maximum likelihood techniques [11]. Therefore the main challenge of forming a funnel multilayer perceptron (MLP) lies in parameterising the weight matrix of an MLP such that it can be used with exact maximum likelihood estimation.

The weight matrix $\mathbf{W} : \mathcal{X}^n \rightarrow \mathcal{Z}^{n-m}$ of a dimension reducing MLP ($0 < m < n$) can be parameterized by an invertible block matrix $\mathbf{R}_{m,m}$ and a generic matrix $\mathbf{W}'_{m,n-m}$ such that

$$\mathbf{W}\mathbf{x} = \mathbf{R}\mathbf{x}_{m:n} + \mathbf{W}'\mathbf{x}_{1:m-1},$$

and the likelihood contribution of the transformation becomes

$$\mathcal{V}(\mathbf{x}) = \log \left(p_\phi \left(\mathbf{x}_{1:m-1} | \mathbf{W}\mathbf{x} + \mathbf{b} \right) \right) + \log (|\det(\mathbf{R})|),$$

where \mathbf{b} is the bias of the layer. Here the matrix \mathbf{R} can have a relatively complex determinant as it only needs to be calculated once per batch update. Also it does not need to be easily invertible, as a lower bound can be found for the likelihood by violating the right inverse condition and instead optimising

$$\mathcal{V}(\mathbf{x}) = \log \left(p_\phi \left(\mathbf{x} | \mathbf{W}\mathbf{x} + \mathbf{b} \right) \right) + \log (|\det(\mathbf{R})|).$$

This transformation also treats some coordinates differently by imposing that the submatrix \mathbf{R} is invertible. This can be partially alleviated by placing flow layers or a fixed random permutation before applying the weight matrix to ensure no specific coordinates are selected to be treated differently. Using results from the existing literature the matrix \mathbf{R} can be parameterized with the LU decomposition [15] or the QR decomposition [16]. Dimension increasing and preserving MLPs are discussed in A.7.2.

4 Experiments

Synthetic Data. We demonstrate the method works qualitatively in Fig. 2 where a funnel layer is used to extend a neural spline flow (NSF) [4] to map from a complex two dimensional distribution to a simple one dimensional distribution (A.5). The matching of the encoded data to the base distribution can be seen on the left, and a faithful distribution of model samples in the middle and the likelihood

on the right. The success of the funnel model on this highly correlated dataset demonstrates its efficacy in handling complex data.

Tabular Data. We next evaluate the performance on selected UCI datasets [17] and the BSDS300 collection of natural images [18] with the same preprocessing and experimental setup as used in masked autoregressive flows [5], as well as utilising the same data and preprocessing [19]. The extensions made to the masked autoregressive flows in Ref. [4] are applied to the models here. We replicate the experimental setup of NSF models [4] as baselines, and extend all models to become funnels as outlined below. No comparisons to SurVAE layers are considered as there are no models that offer comparable performance except for the slice surjection, which we consider to be a simplified instance of a funnel layer where F is the identity (A.3). A baseline VAE model constructed of multilayer perceptrons (MLPs) is also included where the decoder outputs both a mean and standard deviation such that the evidence lower bound (ELBO) [7] can be used as a bound for the likelihood (A.4).

Table 1: Test log likelihood (in nats, higher is better) for UCI datasets and BSDS300, with error bars corresponding to two standard deviations. All models are coupling flows. The F prefix signifies the model has been converted into a funnel with a 25% smaller latent space.

MODEL	POWER	GAS	HEPMASS	MINBOONE	BSDS300
GLOW	0.38 \pm 0.01	12.02 \pm 0.02	-17.22 \pm 0.02	-10.65 \pm 0.45	156.96 \pm 0.28
F-GLOW	0.38 \pm 0.01	12.03 \pm 0.02	-18.25 \pm 0.02	-10.58 \pm 0.44	155.85 \pm 0.85
NSF	0.63 \pm 0.01	13.02 \pm 0.02	-14.92 \pm 0.02	-9.58 \pm 0.48	157.61 \pm 0.28
F-NSF	0.62 \pm 0.02	12.88 \pm 0.02	-15.07 \pm 0.02	-9.50 \pm 0.48	156.90 \pm 0.28
VAE	0.10 \pm 0.02	0.02 \pm 0.05	-22.00 \pm 0.02	-15.39 \pm 0.49	143.70 \pm 0.29
F-MLP	0.36 \pm 0.01	7.94 \pm 0.03	-17.80 \pm 0.02	-12.81 \pm 0.55	153.26 \pm 0.28

The results of the different models are shown in Table 1. In all tables bold numbers indicate the best model on each measure. The funnel models demonstrate competitive performance on all datasets while having a latent space that is 75% the size of the input data dimension. The dimensionality is reduced by replacing one flow layer with a funnel (A.6). For the BSDS300 dataset a diagonal Gaussian ansatz is used to parameterize $p(\mathbf{x}_{1:m-1}|\mathbf{z})$ and for all other datasets a flow is used. The former ansatz results in a slight speed up in the amount of time per training step, and the latter in an increase of up to 20% in time per training step over the baseline equivalents. Notably the VAE likelihood estimate provided by the ELBO shows significantly degraded performance for the same compression as the funnels.

A simple funnel MLP (F-MLP) with a spline activation function performs much better at the density estimation tasks in Table. 1 than a VAE while taking a similar amount of time to train per training step, which is up to six times faster than the flows considered here. These layers can be dramatically improved by combining them with simple flows and selecting the model architecture more carefully, without sacrificing training speed. An F-MLP layer offers a way to build fast performant models for density estimation tasks.

These results show that funnels can be used to model complex datasets from different domains while reducing the size of the latent space with respect to standard flows. The reduced latent space is useful for down stream tasks, as demonstrated in the next section on image generation and anomaly detection tasks.

Image Data. We benchmark funnel models on the CIFAR-10 [20] and downsampled ImageNet64 [21, 22] image datasets with original 8-bit and reduced 5-bit colour depth. In Table 2 we demonstrate that funnel models with a similar number of parameters have similar performance to the benchmark NSF [4] and baseline flow with affine transformations, while having a latent space dimension one sixteenth of the size. We also consider a VAE with a ResNet18 [23] style encoder and decoder with a learnable standard deviation for each channel at the output of the decoder.

There are many ways to build funnel layers for images; here we use standard convolutional kernels (A.8) where the kernel stride is equivalent to the width, which together with the MaxPool layer [11] can be used to turn convolutional neural networks into funnels. The funnel layers here have a diagonal Gaussian ansatz for $p(\mathbf{x}_{1:m-1}|\mathbf{z})$, which is parameterized by a neural network. Each tile of

the convolution is reconstructed separately, with an amortized ansatz for all tiles. Samples drawn from funnel models are shown in Fig. 3.

Table 2: Test-set bits per dimension (BPD) and latent space size (LS) for CIFAR-10 and ImageNet64 models, lower is better for both measures. The F prefix signifies the model has been converted into a convolutional funnel model. *Results taken from van den Oord et al. [24] and Kingma et. al [1] respectively.

MODEL	CIFAR-10 5-BIT		CIFAR-10 8-BIT		IMAGENET64 5-BIT		IMAGENET64 8-BIT	
	BPD	LS	BPD	LS	BPD	LS	BPD	LS
BASELINE	1.70	3,072	3.41	3,072	1.81	12,288	3.91	12,288
NSF	1.70	3,072	3.38	3,072	1.77	12,288	3.82	12,288
F-NSF	1.71	192	3.46	192	1.80	768	3.90	768
VAE	4.09	192	7.45	192	3.72	768	6.67	768
GLOW*	1.67	3,072	3.35	3,072	1.76	12,288	3.81	12,288




Figure 3: Samples drawn from a convolutional funnel model for CIFAR-10 5-bit left and ImageNet64 5-bit right.

The base distribution (p_θ in Eq. (1)) is chosen to be a standard normal distribution, and we remark that in all cases samples drawn from closer to the centre of the distribution have more structure in funnel models than their baseline equivalents, suggesting the encoded distribution more closely matches the base distribution as expected given the shape of the typical set of high dimensional normal distributions. In Fig. 4 we show samples from a trained NSF model and a funnel model where samples are generated from a normal distribution at different ‘temperatures’ (T) where samples are generated by sampling from a modified prior $z \sim T p_\theta(z)$. We observe that in funnel models samples drawn at lower temperatures are much more structured than those drawn from equivalent lowered temperatures in the NSF model.

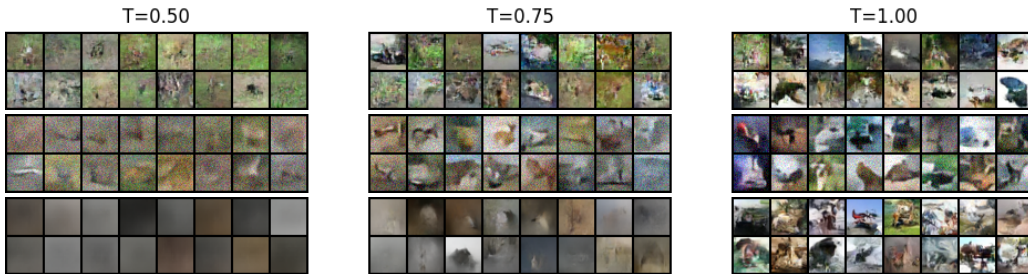


Figure 4: Samples drawn from a funnel model with a latent space of 48 dimensions and BPD of 3.7 (top), a VAE with latent dimension 48 (middle), and an NSF model (bottom) for CIFAR-10 5-bit.

Using standard image quality scores we identify that samples drawn from a funnel model have better visual quality than the baseline equivalent, even though the likelihood score is lower. This demonstrates that funnel models perform better at downstream tasks such as image generation. This can further be interpreted as better matching of the encoded data to the prior.

Anomaly Detection. Out of the box flows often perform poorly on anomaly detection tasks, particularly on images [26]. While flows do work well in some settings [27], this trend is problematic. By

Table 3: Inception score and FID for CIFAR-10 5-BIT for a GLOW based NSF model, a F-NSF model and a ResNet VAE. *Results taken from Ostrovski et. al [25].

MODEL	INCEPTION \uparrow	FID \downarrow
DCGAN*	6.4	37.1
WGAN-GP*	6.5	36.4
PIXELCNN*	4.60	65.93
PIXELQN*	5.29	49.46
NSF	4.15	47.48
F-NSF	4.95	45.52
VAE	1.00	745

Table 4: Bits per dimension scores for models trained on MNIST evaluated on both MNIST and fashion MNIST and the ratio of the likelihoods on both datasets.

MODEL	LATENT SIZE	MNIST \downarrow	F-MNIST \uparrow	RATIO \uparrow
VAE	4	7.30	10.40	1.42
	16	6.95	12.66	1.82
F-NSF	4	1.21	23.42	19.36
	16	1.28	11.32	8.84
F-MLP	4	2.86	30.98	10.83
	16	2.84	28.59	10.07
NSF	768	1.10	4.80	4.36

training models on the MNIST dataset of handwritten digits [28] and evaluating the likelihood scores on the fMNIST dataset of fashion products [29] we identify that funnel models significantly improve the difference in the likelihood scores, as shown in Table. 4. Here we again use the ELBO to provide an out of the box anomaly metric for VAEs. This test demonstrates that funnel models are again more useful for downstream tasks than standard out of the box flows and VAEs. We further remark that a simple F-MLP with spline activations performs significantly better than the VAE model, both of which are constructed from MLPs.

5 Conclusions and discussion

We have developed the funnel layer, a method that allows flows to be extended to have lower dimensional latent spaces while maintaining exact likelihood estimation and similar or better performance with a similar number of parameters. The reduced latent space of a funnel model is shown to be useful for downstream tasks such as generation and anomaly detection and helps with scalability. Crucially the funnel layer allows new transformations to be constructed and could improve the use of exact likelihood methods on tasks that require fast sampling with high dimensional data [10, 30, 31].

Acknowledgments and Disclosure of Funding

The authors would like to acknowledge funding through the SNSF Sinergia grant called Robust Deep Density Models for High-Energy Particle Physics and Solar Flare Analysis (RODEM) with funding number CRSII5_193716. The authors would like to thank Joakim Tutt for useful discussions.

References

- [1] D. P. Kingma and P. Dhariwal, “Glow: Generative flow with invertible 1x1 convolutions,” *arXiv preprint arXiv:1807.03039*, 2018.
- [2] E. G. Tabak and C. V. Turner, “A family of nonparametric density estimation algorithms,” *Communications on Pure and Applied Mathematics*, vol. 66, no. 2, pp. 145–164, 2013.
- [3] L. Dinh, J. Sohl-Dickstein, and S. Bengio, *Density estimation using real nvp*, 2017. arXiv: 1605.08803 [cs.LG].

- [4] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, *Neural spline flows*, 2019. arXiv: 1906.04032 [stat.ML].
- [5] G. Papamakarios, T. Pavlakou, and I. Murray, “Masked autoregressive flow for density estimation,” *arXiv preprint arXiv:1705.07057*, 2017.
- [6] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, *Normalizing flows for probabilistic modeling and inference*, 2021. arXiv: 1912.02762 [stat.ML].
- [7] D. P. Kingma and M. Welling, *Auto-encoding variational bayes*, 2014. arXiv: 1312.6114 [stat.ML].
- [8] D. J. Rezende, S. Mohamed, and D. Wierstra, *Stochastic backpropagation and approximate inference in deep generative models*, 2014. arXiv: 1401.4082 [stat.ML].
- [9] E. Nalisnick, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan, *Do deep generative models know what they don’t know?* 2019. arXiv: 1810.09136 [stat.ML].
- [10] C. Krause and D. Shih, *Caloflow: Fast and accurate generation of calorimeter showers with normalizing flows*, 2021. arXiv: 2106.05285 [physics.ins-det].
- [11] D. Nielsen, P. Jaini, E. Hoogetboom, O. Winther, and M. Welling, “Survae flows: Surjections to bridge the gap between vaes and flows,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [12] C.-W. Huang, L. Dinh, and A. Courville, *Augmented normalizing flows: Bridging the gap between generative flows and latent variable models*, 2020. arXiv: 2002.07101 [cs.LG].
- [13] J. Chen, C. Lu, B. Chenli, J. Zhu, and T. Tian, *Vflow: More expressive generative flows with variational data augmentation*, 2020. arXiv: 2002.09741 [stat.ML].
- [14] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung, “Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit,” *Nature*, vol. 405, no. 6789, pp. 947–951, 2000.
- [15] J. B. Oliva, A. Dubey, M. Zaheer, B. Póczos, R. Salakhutdinov, E. P. Xing, and J. Schneider, *Transformation autoregressive networks*, 2018. arXiv: 1801.09819 [stat.ML].
- [16] E. Hoogetboom, R. van den Berg, and M. Welling, *Emerging convolutions for generative normalizing flows*, 2019. arXiv: 1901.11137 [cs.LG].
- [17] D. Dua and C. Graff, *UCI machine learning repository*, 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>.
- [18] D. Martin, C. Fowlkes, D. Tal, and J. Malik, “A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics,” in *Proc. 8th Int’l Conf. Computer Vision*, vol. 2, Jul. 2001, pp. 416–423.
- [19] G. Papamakarios, *Preprocessed datasets for maf experiments*, Zenodo, Jan. 2018. DOI: 10.5281/zenodo.1161203. [Online]. Available: <https://doi.org/10.5281/zenodo.1161203>.
- [20] A. Krizhevsky, “Learning multiple layers of features from tiny images,” Tech. Rep., 2009.
- [21] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, *Imagenet large scale visual recognition challenge*, 2015. arXiv: 1409.0575 [cs.CV].
- [22] Z. M. Ziegler and A. M. Rush, *Latent normalizing flows for discrete sequences*, 2019. arXiv: 1901.10548 [stat.ML].
- [23] K. He, X. Zhang, S. Ren, and J. Sun, *Deep residual learning for image recognition*, 2015. arXiv: 1512.03385 [cs.CV].
- [24] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu, *Pixel recurrent neural networks*, 2016. arXiv: 1601.06759 [cs.CV].
- [25] G. Ostrovski, W. Dabney, and R. Munos, “Autoregressive quantile networks for generative modeling,” in *International Conference on Machine Learning*, PMLR, 2018, pp. 3936–3945.
- [26] P. Kirichenko, P. Izmailov, and A. G. Wilson, *Why normalizing flows fail to detect out-of-distribution data*, 2020. arXiv: 2006.08545 [stat.ML].

- [27] T. Aarrestad, M. van Beekveld, M. Bona, A. Boveia, S. Caron, J. Davies, A. D. Simone, C. Doglioni, J. M. Duarte, A. Farbin, H. Gupta, L. Hendriks, L. Heinrich, J. Howarth, P. Jawahar, A. Jueid, J. Lastow, A. Leinweber, J. Mamuzic, E. Merényi, A. Morandini, P. Moskvitina, C. Nellist, J. Ngadiuba, B. Ostdiek, M. Pierini, B. Ravina, R. R. de Austri, S. Sekmen, M. Touranakou, M. Vaškevičiūtė, R. Vilalta, J. R. Vlimant, R. Verheyen, M. White, E. Wulff, E. Wallin, K. A. Wozniak, and Z. Zhang, *The dark machines anomaly score challenge: Benchmark data and model independent event classification for the large hadron collider*, 2021. arXiv: 2105.14027 [hep-ph].
- [28] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998. DOI: 10.1109/5.726791.
- [29] H. Xiao, K. Rasul, and R. Vollgraf, *Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms*, 2017. arXiv: 1708.07747 [cs.LG].
- [30] J. Alsing, B. Wandelt, and S. Feeney, “Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology,” *Monthly Notices of the Royal Astronomical Society*, vol. 477, no. 3, pp. 2874–2885, Mar. 2018, ISSN: 1365-2966. DOI: 10.1093/mnras/sty819. [Online]. Available: <http://dx.doi.org/10.1093/mnras/sty819>.
- [31] A. Rouhiainen, U. Giri, and M. Münchmeyer, “Normalizing flows for random fields in cosmology,” May 2021. arXiv: 2105.12024 [astro-ph.CO].
- [32] S. Ioffe and C. Szegedy, *Batch normalization: Accelerating deep network training by reducing internal covariate shift*, 2015. arXiv: 1502.03167 [cs.LG].
- [33] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [34] I. Loshchilov and F. Hutter, *Sgdr: Stochastic gradient descent with warm restarts*, 2017. arXiv: 1608.03983 [cs.LG].

A Appendix

A.1 Definitions and notation

For the bound looseness term to be zero the right inverse condition must be satisfied. The right inverse condition can be stated as $f(\mathbf{x}) = \mathbf{z}$ for any $\mathbf{x} \sim p(\mathbf{x}|\mathbf{z})$ and $\mathbf{z} \in \mathbb{Z}$, where \mathbb{Z} is the support of $q(\mathbf{z}) = \mathbb{E}_{p_D(\mathbf{x})}[q(\mathbf{z}|\mathbf{x})]$. Funnel models satisfy this condition because

$$\begin{aligned} p(\mathbf{x}_{m:n}|\mathbf{z}) = \delta(\mathbf{x}_{m:n} - F(\mathbf{z})) &\implies F(\mathbf{z}) = \mathbf{x}_{m:n} \sim p(\mathbf{x}_{m:n}|\mathbf{z}) \\ &\implies F^{-1}(\mathbf{x}_{m:n}) = \mathbf{z}, \forall \mathbf{x}_{m:n} \sim p(\mathbf{x}_{m:n}|\mathbf{z}). \end{aligned}$$

Note also that a diffeomorphism F parameterized by a function $\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})$ is defined such that its inverse can be found given \mathbf{z} using $\theta(\mathbf{x}_{1:m-1}, \mathbf{z})$. This is done by using coupling or autoregressive transformations on $\mathbf{x}_{m:n}$.

A.2 Generative funnel layer

We can also develop a generative funnel surjection. The likelihood contribution for a generative surjection is given by

$$\mathcal{V}(\mathbf{x}, \mathbf{z}) = \lim_{p(\mathbf{z}|\mathbf{x}) \rightarrow \delta(\mathbf{z} - f^{-1}(\mathbf{x}))} \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \frac{p(\mathbf{x}|\mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \right].$$

For a generative funnel transformation we split $\mathbf{z} = \{\mathbf{z}_{1:n}, \mathbf{z}_{n+1:m}\}$ and use a diffeomorphism G such that $\mathbf{z}_{1:n} = G_{\theta(\mathbf{x}, \mathbf{z}_{n+1:m})}^{-1}(\mathbf{x})$ with $\mathbf{z}_{n+1:m} \sim p(\mathbf{z}_{n+1:m}|\mathbf{x})$. Then following the same steps as for the inference surjection we can use

$$\begin{aligned} p(\mathbf{x}|\mathbf{z}) &= \delta(\mathbf{x} - G_{\theta(\mathbf{z}_{1:n}, \mathbf{z}_{n+1:m})}(\mathbf{z}_{1:n})) \\ &= \delta(\mathbf{z}_{1:n} - G_{\theta(\mathbf{z}_{1:n}, \mathbf{z}_{n+1:m})}^{-1}(\mathbf{x})) |\det(J_{\theta}(\mathbf{z}))|, \end{aligned}$$

to write

$$\begin{aligned} \mathcal{V}(\mathbf{x}) &= \int d\mathbf{z} q(\mathbf{z}|\mathbf{x}) \log \left(\frac{\delta(\mathbf{z}_{1:n} - G_{\theta(\mathbf{z}_{1:n}, \mathbf{z}_{n+1:m})}^{-1}(\mathbf{x})) |\det(J_{\theta}(\mathbf{z}))|}{q(\mathbf{z}|\mathbf{x})} \right) \\ &= \int d\mathbf{z}_{n+1:m} q(\mathbf{z}_{n+1:m}|\mathbf{x}) \log \left(\frac{|\det(J_{\theta}(\mathbf{z}))|}{q(\mathbf{z}_{n+1:m}|\mathbf{x})} \right), \end{aligned} \tag{9}$$

with

$$J_{\theta}(\mathbf{z}) = \left. \frac{\partial G_{\theta(\mathbf{z}_{1:n}, \mathbf{z}_{n+1:m})}^{-1}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x} = G_{\theta(\mathbf{z}_{1:n}, \mathbf{z}_{n+1:m})}(\mathbf{z}_{1:n})}.$$

A.3 Relationship of funnel to slice surjection

The funnel layer is related to the slice surjection presented in Ref. [11]. There are several key differences between the slice surjection and the funnel layer, but before discussing them we detail how the slice surjection can be related to the funnel when composing with invertible neural networks (INNs).

Define a coupling INN F_{θ_s} that acts on \mathbf{x}

$$F_{\theta_s}(\mathbf{x}) = (\mathbf{x}_{1:m-1}, F_{\theta_s(\mathbf{x}_{1:m-1})}(\mathbf{x}_{m:n})), \tag{10}$$

and a likelihood contribution of

$$\mathcal{V}(\mathbf{x}) = \log(|\det J_{F_{\theta_s}}(\mathbf{x}_{m:n})|). \tag{11}$$

Combining this with a slice surjection S that acts as

$$S(F_{\theta_s}(\mathbf{x})) = F_{\theta_s(\mathbf{x}_{1:m-1})}(\mathbf{x}_{m:n}), \tag{12}$$

the total likelihood contribution will be

$$\mathcal{V}(\mathbf{x}) = \log \left(p \left(\mathbf{x}_{1:m-1} \middle| F_{\theta_s(\mathbf{x}_{1:m-1})}(\mathbf{x}_{m:n}) \right) \right) + \log(|\det J_{F_{\theta_s}}(\mathbf{x}_{m:n})|). \quad (13)$$

This is a special case of the likelihood contribution in Eq. (6), and it requires selecting a particular invertible neural network in Eq. (10). The funnel layer provides a unified perspective on this combination and allows for more flexibility in terms of how the dependence on $\mathbf{x}_{1:m-1}$ is propagated as well as the choice of the function $F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})$ that is used in the transformation.

The slice layer is restricted to being composed with INNs for the function F_{θ_s} . This is not the case for the funnel method, and this freedom allows convolutional kernels to be used with maximum likelihood estimation within the funnel framework as we show in Appendix A.7.1. In general the funnel method requires that $F_{\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})}^{-1}(\mathbf{x}_{m:n})$ is invertible for $\mathbf{x}_{m:n}$ given $\mathbf{x}_{1:m-1}$ and has a tractable determinant, this includes functions that are not typically considered as INNs.

When the INN in Eq. (10) is parameterized in terms of a two layer neural network $\theta = \theta_1 \circ \theta_0$ that acts on $(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})$ the dependence on $\mathbf{x}_{1:m-1}$ is propagated as

$$\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n}) = \theta_1(\theta_0(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})), \quad (14)$$

which is the standard way of building a coupling INN. In comparison, when using the funnel method the above would be possible, but it is also natural to consider building the dependence as

$$\theta(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n}) = \theta_1(\mathbf{x}_{1:m-1}, \theta_0(\mathbf{x}_{1:m-1}, \mathbf{x}_{m:n})), \quad (15)$$

where the resulting transformation has an increased dependence on $\mathbf{x}_{1:m-1}$. In comparison the INN in Eq. (10) $\mathbf{x}_{1:m-1}$ would normally be transformed in the next layer of the model, and so it would be unnatural to propagate the dependence on $\mathbf{x}_{1:m-1}$ in the way shown in Eq. 15.

Another benefit of the funnel method is clarity of thinking about the different layers. By building the generic composition directly the structure of the model can be viewed holistically.

The slice surjection can be viewed as a particular case of the funnel where the forward transformation F is the identity. Further, any dependence that can be built into a slice surjection through composition with INNs can also be built into a funnel layer using the same composition.

The inference version of the slice surjection was related to the multi-scale architecture from Ref. [3]. Here we demonstrate that its generalisation, the funnel, can be used to build dimension reducing transformations which are effective on a wide variety of datasets.

A.4 VAE ELBO

The evidence lower bound (ELBO) of a variational autoencoder (VAE) is given by

$$\log p(\mathbf{x}) = \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}|\mathbf{z})] - \mathbb{D}_{\text{KL}}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]. \quad (16)$$

In this work we calculate the KL divergence term with the reparameterisation trick as is standard, and the expectation value is estimated with a single sample and we assume a gaussian distribution for the reconstruction using a neural network $g_\theta(\mathbf{z}) = [\mu(\mathbf{z}), \sigma(\mathbf{z})]$ such that

$$p(\mathbf{x}|\mathbf{z}) = \frac{1}{\sigma(\mathbf{z})\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\mathbf{x}-\mu(\mathbf{z})}{\sigma(\mathbf{z})}\right)^2}. \quad (17)$$

A.5 Plane data distribution

The dataset that is modelled in Fig. 2 is called four circles and the data distribution is shown in Fig. 5.

On this dataset a neural spline flow (NSF) [4] with the last layer replaced by a funnel was used. The NSF was formed from two autoregressive rational quadratic layers with 50 bins in each layer, a tail bound of 4 and a standard normal as the base distribution. The parameters of the splines are learned with two residual blocks composed of two multilayer perceptrons (MLP) with 256 hidden units. The inverse density $p_\phi(\mathbf{x}_1|\mathbf{z})$ is composed of an NSF with the same structure but with 64 nodes in the MLP and ten bins per layer.




Figure 5: The data distribution (Four circles) used to train the model in Fig. 2.

A.6 UCI experimental details

All parameters are chosen to replicate the experiments from Ref. [4] and funnel models are built such that they have less parameters than their non-funnel baselines. No model selection has been performed for the funnel approaches, except tuning to match the numbers of parameters used. The norm of the gradients is clipped to the range $[-5, 5]$. Hyperparameter settings are shown for coupling flows in Table 5. The dimensionality and number of training data points in each dataset is included for reference. When using a flow to parameterize the density $p(\mathbf{x}_{1:m-1}|\mathbf{z})$ the same type of flow as is used in the modelling is used, with half the number of flow steps and one third as many hidden features in the residual blocks. For BSDS300 a Gaussian ansatz was used for this density, with an MLP composed of three layer with 128 nodes and ReLU activations [14] to predict the mean and diagonal covariance matrix of the Gaussian.

Table 5: The hyperparameters used for uci datasets and BSDS300 for both flows and MLPs.

	POWER	GAS	HEPMASS	MINBOONE	BSDS300
DIMENSION	6	8	21	43	63
DATA POINTS	1,615,917	852,174	315,123	29,556	1,000,000
BATCH SIZE	512	512	256	128	512
TRAINING STEPS	400,000	400,000	400,000	200,000	400,000
LEARNING RATE	0.0005	0.0005	0.0005	0.0003	0.0005
FLOW STEPS	10	10	20	10	10
RESIDUAL BLOCKS	2	2	1	1	1
HIDDEN FEATURES	256	256	128	32	128
BINS	8	8	8	4	8
DROPOUT	0.0	0.1	0.2	0.2	0.2
DIM. REDUCTION	2	2	5	10	15
FUNNEL LEVEL	2	2	0	4	4
VAE					
MLP WIDTH	256	64	256	512	512
MLP DEPTH	2	8	2	2	2
LEARNING RATE	0.003	0.005	0.005	0.001	0.003

For the VAEs used in these experiments a hyper parameter scan was performed to find the best model, with the number of parameters varying across the scan but allowing for both more and less parameters than the flow models. In all cases the DROPOUT= 0 and the BATCH SIZE was the same as that used for flows. Both batch normalization and layer normalization were included in the scan but never resulted in the best performance. The parameters in Table. 6 were searched to find the best VAE for each dataset.

Each layer of the F-MLP model used a fixed random permutation followed by an F-MLP layer where an invertible matrix parameterized by the LU decomposition was used for \mathbf{R} . For dimension preserving layers just the matrix \mathbf{R} was used. A rational quadratic spline with 10 bins and a tail bound of 2 was used as an activation for each layer. Each experiment used two dimension preserving layers, followed by one layer to map to a vector with half the dimension of the input and the latent space dimension, followed by a dimension preserving layer, a layer to map to the latent space, and

Table 6: The parameters searched to find the best VAE for each of the tabular datasets.

MLP WIDTH	64, 128, 256, 512
MLP DEPTH	2, 3, 4, 6, 8
DROPOUT	0.0, 0.2, 0.5
BATCH SIZE	128, 256, 512, 1024
LEARNING RATE	0.0003, 0.0001, 0.00008, 0.00005

then three dimension preserving layers. All other parameters were held the same as those used for the flows in Table. 5 and no hyper parameter scan was performed.

A.7 Standard layers

In this section we detail how both standard convolutional kernels and MLPs can be converted into funnel models.

A.7.1 Convolutional kernel funnels

In this section we outline the method for turning convolutional kernels into a funnel layer. We will recover the simple one dimensional example and then extend to more complex data types.

Given a one dimensional vector $\mathbf{x} \in \mathbb{R}^n$ the forward transformation of the funnel can be factorised to form a 2×1 convolutional kernel $([a, b])$ with stride two as follows

$$q(\mathbf{z}|\mathbf{x}) = \prod_{i=1}^{n/2} q(z_i|x_{2i-1}, x_{2i}) \quad (18)$$

with,

$$\begin{aligned} z_i &= F_{\theta(x_{2i-1}, x_{2i})}^{-1}(x_{2i-1}) \\ &= ax_{2i-1} + bx_{2i} \end{aligned}$$

where $[a, b]$, are the parameters of the convolutional kernel. The Jacobian of each individual kernel is given by

$$J_i = \left[\frac{dF_{\theta(x_{2i-1}, x_{2i})}^{-1}(x_{2i-1})}{dx_{2i-1}} \right] = [a],$$

and the total Jacobian will be given by $\frac{n}{2}a$ such that the total likelihood contribution is given by

$$\mathcal{V}(\mathbf{x}) = \log \left(p_\phi \left(\{x_{2i}\}_{i=1}^{n/2} | \mathbf{z} \right) \right) + \log \left(\frac{n}{2} |a| \right).$$

Given $x'_{2i} \sim p_\phi \left(\{x_{2i}\}_{i=1}^{n/2} | \mathbf{z} \right)$, the kernel $[a, b]$ and \mathbf{z} we can find x_{2i-1}

$$x_{2i-1} = J_i^{-1}(z_i - bx'_{2i}), \quad (19)$$

for all i such that the right inverse condition is satisfied. Here the different elements of the kernel play different roles in the transformation with a appearing directly in the loss and b playing an implicit role in p_ϕ , we also chose x_{2i-1} arbitrarily to play the role of $\mathbf{x}_{m:n}$. As we have chosen a stride of two, this means different features in the input vector will also play different roles. The issue of certain features playing a special role is resolved if we take a stride of one, as then the forward transformation is given by

$$q(z_i|x_i, x_{i+1}) = \delta(z_i - ax_i - bx_{i+1}),$$

for all $i \in [1, n-1]$, and all features are multiplied by the special element of the kernel (a). But, in this setting we have the autoregressive property and so

$$p_\phi(\mathbf{x}|\mathbf{z}) = p_\phi(x_n|\mathbf{z}),$$

because if we have access to x_n (or a sample from $p_\phi(x_n|z)$) then x can be unambiguously recovered using the recursion relation

$$x_{n-i} = J_i^{-1}(z_{n-i} - bx'_{n-i+1}),$$

for all $i \in [1, n-1]$. In settings with low dimensional data this is a very useful property, but for images it makes the sampling process very slow. To speed up sampling the likelihood of the model can be bounded from below by violating the right inverse condition, and estimating $p_\phi(x|z)$ directly. To match the appearance of a in the likelihood contribution it is possibly better to have a generative process p_ϕ that factorises as

$$p_\phi(x|z) = p_\phi(x_1|ax_1 + bx_2) \prod_{i=2}^n p_\phi(x_i|ax_{i-1} + bx_i),$$

where the conditioning in all densities except the first enhances the role played by b . This is the form of the generative process that we use in all studies, were the same model p_ϕ is used in each inverse block. A better approach would be to remove the ambiguity associated with the choice of kernel element to appear in the loss.

The above approach can be extended to three channel images. Given a kernel defined by $a_{i,j}^{k,l}$, where $a \in \mathbb{R}$ and i, j are the indices of each matrix in the kernel, k corresponds to the number of output channels and l to the number of input channels. Fixing $k = l$ and $i = j = 2$ and choosing (again arbitrarily) to take the bottom right component of each matrix to be $x_{m:n}$ the Jacobian is given by

$$J = \begin{bmatrix} a_{2,2}^{1,1} & a_{2,2}^{2,1} & a_{2,2}^{3,1} \\ a_{2,2}^{1,2} & a_{2,2}^{2,2} & a_{2,2}^{3,2} \\ a_{2,2}^{1,3} & a_{2,2}^{2,3} & a_{2,2}^{3,3} \end{bmatrix}. \quad (20)$$

The determinant of the Jacobian has to be calculated once per batch, and so allows for efficient training. The inverse Jacobian has to be calculated once for every trained model, allowing for efficient sampling.

Samples can be generated efficiently using the following steps. Defining the kernel defined by $a_{i,j}^{k,l}$ to act as $F^{-1}(x_{1:m-1}, x_{m:n})$ then for each tile j in the forward convolution we can find the coordinates $x_{m:n}^j$ by sampling $x_{1:m-1}^j$ from $p_\phi(x_{1:m-1}|z_j)$ and using

$$x_{m:n}^j = J^{-1}(z_j - F^{-1}(x_{1:m-1}^j, \mathbf{0}_{\mathbb{R}^3})). \quad (21)$$

to ensure the right inverse condition will be satisfied in the case where the stride of the kernel is equivalent to its width. This procedure can be parallelised across all tiles of a convolution in standard machine learning libraries, leading to an efficient sampling process.

This generalizes to all $\{i, j\}$ but the requirement that $k = l$ is fixed to ensure the matrix J is square. The same considerations apply when using a stride that is less than the width of the kernel.

Other approaches with funnel models can be developed, such as using coupling transformations and dropping some of the untransformed channels after the transformation has been applied. In practise this works but requires more parameters to function and so was not included in the comparisons in Table. 2. This approach could be used, for example, to transform grey scale to color images.

A.7.2 MLP funnels

Dimension increasing funnels cannot be developed within this framework as the inference surjection is needed, but this requires a stochastic $q(z|x)$ which cannot be done with standard MLPs.

A.8 Image experimental details

For the generative modelling of images experiments we use a rational quadratic spline based architecture [4] that is based on a Glow-like model [1]. Our approach exactly replicates all experimental details from Ref. [4], where the multi-scale architecture that is used here can be viewed as a simple slice surjection [11]. In the following we replicate the information provided in Ref. [4]

The architecture consist of multiple steps for each level in a multiscale architecture [3], where each step consists of an actnorm layer, an invertible 1×1 convolution and a coupling transform. For our

RQ-NSF (C) model, rational quadratic coupling transforms parameterized by residual convolutional networks are used, and a 1×1 convolution at the end of each level of transform. For CIFAR-10 experiments dimensions are not factored out at the end of each level, but the squeezing operation is used to trade spatial resolution for depth. All experiments used 3 residual blocks and batch normalization [32] in the residual networks which parameterize the coupling transforms. We use 7 steps per level for all experiments, resulting in a total of 21 coupling transforms for CIFAR-10, and 28 coupling transform for ImageNet64.

We use the Adam optimizer [33] with default β_1 and β_2 values. An initial learning rate of 0.0005 is annealed to 0 following a cosine schedule [34]. We train for 100,000 steps for 5-bit experiments, and for 200,000 steps for 8-bit experiments. Final values for all hyperparameters are reported in Table 7. We use a single Nvidia RTX 3090 GPU card per image experiment. Training for 200,000 steps takes about 4 days on this card. All funnel experiments are run for 12 hours.

As we are interested in dimension reducing transformations, we only used convolutions where the width of the kernel and the stride are equivalent. In all cases we use a convolutional kernel of size 2 with 3 channels and no padding. We apply an actnorm layer, an invertible 1×1 convolution and a coupling transformation (the first block of the NSF architecture) followed by the convolutional kernel and then the rest of the NSF architecture. The Gaussian ansatz, used to parameterize the log likelihood of the tiles in the kernel, is formed from a three layer MLP with 128 nodes in each layer and ReLU activations. In this case the ansatz is trained to invert the tile of each kernel individually, given all neighbours of the corresponding element z_i . This results in a tile of z_i elements of size 3 with 3 channels giving 27 inputs to infer the $(2 \times 2 - 1) \times 3 = 9$ elements of $\mathbf{x}_{1:m-1}$ in each tile in the input image.

Table 7: The hyperparameters used for generative modelling of images.

DATASET	BITS	BATCH SIZE	LEVELS	HIDDEN CHANNELS	BINS	DROPOUT
CIFAR-10	5	512	3	64	2	0.2
	8	512	3	96	4	0.2
IMAGENET64	5	256	4	96	8	0.1
	8	256	4	96	8	0.0

For the VAEs in these experiments a ResNet18 encoder [23] with a single linear layer to map from the output of the ResNet18 model to the latent space. The same convolutions are repeated for the decoder with convolutional kernels for the upsampling.

A.9 Anomaly detection

For the NSF model we use rational quadratic coupling transforms parameterized by residual convolutional networks, and a 1×1 convolution at the end of each level of transform. All experiments use 3 residual blocks and batch normalization [32] and 7 steps per level. Squeezing is applied between different layers to exchange spatial resolution for depth.

For the F-NSF we used a convolution of size (2, 2) and stride two in each layer, with zero padding after the first convolution to ensure the dimension of the image remains even. To compress to 16 dimensions three layers were used, and to compress to 4 dimensions four layers were used. The same layers were used as in the case of the NSF but no squeezing was applied.

The F-MLP model used a fixed random permutation and an invertible matrix parameterized by the LU decomposition was used for \mathbf{R} . A fixed per dimension rational quadratic spline with 5 bins and a tial bound of 4 [4] was used as an activation and the network compressed the input by first flattening the image and then applying five F-MLP layers as $768 \rightarrow 512 \rightarrow 256 \rightarrow 128 \rightarrow 64 \rightarrow \text{LS}$ where the latent space size (LS) is either 4 or 16.

The VAE model was parameterized by linear layers with RELU activations [14], the encoder and decoder had the same structure and the network compressed the image by first flattening and then applying linear layers as $768 \rightarrow 512 \rightarrow 512 \rightarrow 512 \rightarrow \text{LS}$ where LS is either 4 or 16.

Table 8: The number of parameters used in each of the models in Table. 4.

MODEL	LATENT SIZE	PARAM
VAE	4	1,861,401
	16	1,879,857
F-NSF	4	7,286,164
	16	5,464,623
F-MLP	4	1,890,324
	16	1,888,020
NSF	768	4,724,430