

# UC Berkeley

## UC Berkeley Previously Published Works

### Title

Likelihood Contribution based Multi-scale Architecture for Generative Flows

### Permalink

<https://escholarship.org/uc/item/9mv4r0t8>

### Authors

Das, Hari Prasanna

Abbeel, Pieter

Spanos, Costas J

### Publication Date

2019-08-05

Peer reviewed

---

# Likelihood Contribution based Multi-scale Architecture for Generative Flows

---

**Hari Prasanna Das, Pieter Abbeel and Costas J. Spanos**  
Department of Electrical Engineering and Computer Sciences  
University of California, Berkeley  
{hpdas, pabbeel, spanos}@berkeley.edu

## Abstract

Deep generative modeling using flows has gained popularity owing to the tractable exact log-likelihood estimation with efficient training and synthesis process. However, flow models suffer from the challenge of having high dimensional latent space, the same in dimension as the input space. An effective solution to the above challenge as proposed by Dinh et al. (2016) is a multi-scale architecture, which is based on iterative early factorization of a part of the total dimensions at regular intervals. Prior works on generative flow models involving a multi-scale architecture perform the dimension factorization based on static masking. We propose a novel multi-scale architecture that performs data-dependent factorization to decide which dimensions should pass through more flow layers. To facilitate the same, we introduce a heuristic based on the contribution of each dimension to the total log-likelihood which encodes the importance of the dimensions. Our proposed heuristic is readily obtained as part of the flow training process, enabling the versatile implementation of our likelihood contribution based multi-scale architecture for generic flow models. We present such implementations for several state-of-the-art flow models and demonstrate improvements in log-likelihood score and sampling quality on standard image benchmarks. We also conduct ablation studies to compare the proposed method with other options for dimension factorization.

## 1 Introduction

Deep Generative Modeling aims to learn the embedded distributions and representations in input (especially unlabelled) data, requiring no/minimal human labelling effort. The representations learnt can then be utilized in a number of downstream tasks such as semi-supervised learning (Kingma et al., 2014; Odena, 2016), synthetic data augmentation and adversarial training (Cisse et al., 2017), text analysis and model based control etc. The repository of deep generative modeling majorly includes likelihood based models such as autoregressive models, latent variable models, flow based models and implicit models such as generative adversarial networks (GANs). Autoregressive models (Salimans et al., 2017; Oord et al., 2016b,a; Chen et al., 2017) achieve exceptional log-likelihood score on many standard datasets, indicative of their power to model the inherent distribution. But, they suffer from a slow sampling process. Latent variable models such as variational autoencoders (Kingma and Welling, 2013) tend to better capture the global feature representation in data, but do not offer an exact density estimate. Implicit generative models such as GANs optimize a generator and a discriminator in a min-max fashion have recently become popular for their ability to synthesize realistic data (Karras et al., 2018; Engel et al., 2019). But, GANs neither offer a latent space suitable for further downstream tasks, nor do they perform density estimation.

Flow based generative models (Dinh et al., 2016; Kingma and Dhariwal, 2018) perform exact density estimation with fast inference and sampling, due to their parallelizability. They also provide an information rich latent space suitable for many applications. However, the dimension of latent space

for flow based generative models is same as the high-dimensional input space, by virtue of bijective nature of flows. This poses a bottleneck for flow models to scale with increasing input dimensions due to computational complexity. An effective solution to the above problem is a multi-scale architecture, introduced by Dinh et al. (2016), which performs iterative early gaussianization of a part of the total dimensions at regular intervals of flow layers. This not only makes the model computational and memory efficient but also aids in distributing the loss function throughout the network for better training. Many prior works including Kingma and Dhariwal (2018); Behrman et al. (2018); Atanov et al. (2019); Kumar et al. (2019) implement multi-scale architecture in their flow models, but use static masking methods for factorization of dimensions.

We propose a multi-scale architecture which performs data dependent factorization to decide which dimensions should pass through more flow layers. For the decision making, we introduce a heuristic based on the amount of log-likelihood contribution by each dimension, which in turn signifies their individual importance. Since in the proposed architecture, the heuristic is obtained as part of the flow training process, it can be universally applied to generic flow models. We present such implementations for flow models based on affine/additive coupling as well as ordinary differential equations (ODE) and achieve quantitative and qualitative improvements. We also perform ablation studies to establish the novelty of our method. Summing up, the contributions of our research are,

**Contributions:**

1. A log-determinant based heuristic which entails the contribution by each dimension towards the total log-likelihood in a multi-scale architecture.
2. A multi-scale architecture based on the above heuristic performing data-dependent splitting of dimensions, implemented for several classes of flow models.
3. Quantitative and qualitative analysis of above implementations and an ablation study

To the best of our knowledge, we are the first to propose a data-dependent splitting of dimensions in a multi-scale architecture.

## 2 Background

### 2.1 Flow-based Generative Models

Let  $\mathbf{x}$  be a high-dimensional random vector with unknown true distribution  $p(\mathbf{x})$ . The following formulation is directly applicable to continuous data, and with some pre-processing steps such as dequantization (Uria et al., 2013; Salimans et al., 2017; Ho et al., 2019) to discrete data. Let  $\mathbf{z}$  be the latent variable with a known standard distribution  $p(\mathbf{z})$ , such as a standard multivariate gaussian. Using an i.i.d. dataset  $\mathcal{D}$ , the target is to model  $p_{\theta}(\mathbf{x})$  with parameters  $\theta$ . A flow,  $\mathbf{f}_{\theta}$  is defined to be an invertible transformation that maps observed data  $\mathbf{x}$  to the latent variable  $\mathbf{z}$ . A flow is invertible, so the inverse function  $\mathcal{T}$  maps  $\mathbf{z}$  to  $\mathbf{x}$ , i.e.

$$\mathbf{z} = \mathbf{f}_{\theta}(\mathbf{x}) = \mathcal{T}^{-1}(\mathbf{x}) \text{ and } \mathbf{x} = \mathcal{T}(\mathbf{z}) = \mathbf{f}_{\theta}^{-1}(\mathbf{z}) \tag{1}$$

The log-likelihood can be expressed as,

$$\log p_{\theta}(\mathbf{x}) = \log p(\mathbf{z}) + \log \left| \det \left( \frac{\partial \mathbf{f}_{\theta}(\mathbf{x})}{\partial \mathbf{x}^T} \right) \right| \tag{2}$$

where  $\frac{\partial \mathbf{f}_{\theta}(\mathbf{x})}{\partial \mathbf{x}^T}$  is the Jacobian of  $\mathbf{f}_{\theta}$  at  $\mathbf{x}$ . The invertible nature of a flow allows it to be capable of being composed of other flows of compatible dimensions. In practice, flows are constructed by composing a series of component flows. Let the flow  $\mathbf{f}_{\theta}$  be composed of  $K$  component flows, i.e.  $\mathbf{f}_{\theta} = \mathbf{f}_{\theta_K} \circ \mathbf{f}_{\theta_{K-1}} \circ \dots \circ \mathbf{f}_{\theta_1}$  and the intermediate variables be denoted by  $\mathbf{y}_K, \mathbf{y}_{K-1}, \dots, \mathbf{y}_0 = \mathbf{x}$ . Then the log-likelihood of the composed flow is,

$$\log p_{\theta}(\mathbf{x}) = \underbrace{\log p(\mathbf{z})}_{\text{Log-latent density}} + \log \left| \det \left( \frac{\partial (\mathbf{f}_{\theta_K} \circ \mathbf{f}_{\theta_{K-1}} \circ \dots \circ \mathbf{f}_{\theta_1})(\mathbf{x})}{\partial \mathbf{x}^T} \right) \right| \tag{3}$$

$= \sum_{i=1}^K \log |\det(\partial \mathbf{y}_i / \partial \mathbf{y}_{i-1}^T)|$  (Log-det)

which follows from the fact that  $\det(A \cdot B) = \det(A) \cdot \det(B)$ . In our work, we refer the first term in Equation 3 as *log-latent-density* and the second term as *log-determinant (log-det)*. The reverse path, from  $\mathbf{z}$  to  $\mathbf{x}$  can be written as a composition of inverse flows,  $\mathbf{x} = \mathbf{f}_{\theta}^{-1}(\mathbf{z}) = \mathbf{f}_{\theta_1}^{-1} \circ \mathbf{f}_{\theta_2}^{-1} \circ \dots \circ \mathbf{f}_{\theta_K}^{-1}(\mathbf{z})$ . Confirming with above properties for a flow, different types of flows can be constructed (Kingma and Dhariwal, 2018; Dinh et al., 2016, 2014; Behrman et al., 2018; Chen et al., 2019).

## 2.2 Multi-scale Architecture

Multi-scale architecture is a design choice for latent space dimensionality reduction of flow models, in which part of the dimensions are factored out/early gaussianized at regular intervals, and the other part is exposed to more flow layers. The process is called dimension factorization. In the problem setting as introduced in Section 2.1, the factoring operation can be mathematically expressed as,

$$\begin{aligned} \mathbf{y}_0 = \mathbf{x}, (\mathbf{z}_{l+1}, \mathbf{y}_{l+1}) &= \mathbf{f}_{\theta_{l+1}}(\mathbf{y}_l), \quad l \in \{0, 1, \dots, K-2\} \\ \mathbf{z}_K &= \mathbf{f}_{\theta_K}(\mathbf{y}_{K-1}), \mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_K) \end{aligned}$$

The factoring of dimensions at early layers has the benefit of distributing the loss function throughout the network (Dinh et al., 2016) and optimizing the amount of computation used by the model.

## 3 Likelihood Contribution based Multi-scale Architecture

In a multi-scale architecture, it is apparent that the network will better learn the distribution of dimensions getting exposed to more layers of flow as compared to the ones which get factored at a finer scale (earlier layer). The method of dimension splitting proposed by prior works such as (Dinh et al., 2016; Kingma and Dhariwal, 2018; Behrmann et al., 2018) are static in nature and do not distinguish between importance of different dimensions. In this section, we introduce the general framework for likelihood contribution based heuristic and associated multi-scale architecture along with its integration with flow training process.

### 3.1 Likelihood Contribution based Heuristic

Recall from Equation 3 that the log-likelihood is composed of two terms, the log-latent density term and the log-det term. We focus on the log-det term since it depends on the modeling of flow layers.

Let the dimension of the input (images in our case) space  $\mathbf{x}$  be  $s \times s \times c$ , where  $s$  is the height/width and  $c$  is the number of channels. For the following formulation, let us assume NO dimensions were gaussianized early so that we have access to log-det term for all dimensions at each flow layer, and the dimension at all intermediate layers remain the same (i.e.  $s \times s \times c$ ). We apply a flow ( $\mathbf{f}_\theta$ ) with  $K$  component flows to  $\mathbf{x}, \mathbf{z}$  pair, so that  $\mathbf{z} = \mathbf{f}_\theta(\mathbf{x}) = \mathbf{f}_{\theta_K} \circ \mathbf{f}_{\theta_{K-1}} \circ \dots \circ \mathbf{f}_{\theta_1}(\mathbf{x})$ . The intermediate variables are denoted by  $\mathbf{y}_K, \mathbf{y}_{K-1}, \dots, \mathbf{y}_0$  with  $\mathbf{y}_K = \mathbf{z}$  (since no early gaussianization was performed) and  $\mathbf{y}_0 = \mathbf{x}$ . The log-det term at layer  $l$ ,  $L_d^{(l)}$ , is given by,

$$[L_d^{(l)}]_{scaler} = \sum_{i=1}^l \log |\det(\partial \mathbf{y}_i / \partial \mathbf{y}_{i-1}^T)| \quad (4)$$

The log-det of the jacobian term encompasses contribution by all the  $s \times s \times c$  dimensions. If we decompose it to obtain the individual contribution by the dimensions (we discuss explicitly on how to perform this decomposition in Sec. 3.2) towards the total log-det ( $\sim$  total log-likelihood). The log-det term can be viewed (with slight abuse of notations) as a  $s \times s \times c$  tensor corresponding to each of the dimensions, summed over the flow layers till  $l$ ,

$$[L_d^{(l)}]_{s \times s \times c} = \sum_{i=1}^l [d_{i-1}^{(\alpha, \beta, \gamma)}]_{s \times s \times c},$$

$$\text{where } \alpha, \beta \in \{0, \dots, s\} \text{ and } \gamma \in \{0, \dots, c\}, \text{ s.t. } \sum_{\alpha, \beta, \gamma} d_{i-1}^{(\alpha, \beta, \gamma)} = \log |\det(\partial \mathbf{y}_i / \partial \mathbf{y}_{i-1}^T)|$$

The entries in  $[L_d^{(l)}]_{s \times s \times c}$  having higher value were scaled up more, and correspond to dimensions which are more sensitive to changes in input, so the flow can learn more by processing them through more layers. So, we can use the *likelihood contribution (in the form of log-det term) by each dimension* as a heuristic for deciding which variables should be gaussianized early in a multi-scale architecture.

### 3.2 Estimation of Per-Dimensional Likelihood Contribution for various types of Flows

The likelihood (log-det) contributions per dimension can be obtained by decomposition of the overall log-det of the jacobian. Now, we describe the decomposition process for various types of flow models. The log-det per dimension after decomposition is averaged across the samples in training set, so as to obtain an overall representative likelihood contribution by each dimension.

### 3.2.1 Affine coupling based flows

*RealNVP (Dinh et al., 2016)*: For RealNVP with affine coupling layers, the logarithm of individual diagonal elements of jacobian, summed over layers till  $l$  provides the per-dimensional likelihood contribution components at layer  $l$ .

*Glow (Kingma and Dhariwal, 2018)*: Unlike RealNVP where the log-det terms for each dimension can be expressed as log of corresponding diagonal element of jacobian, Glow contains  $1 \times 1$  convolution blocks having non-diagonal log-det term. For a  $s \times s \times c$  tensor, the log-det term is  $s \cdot s \cdot \log |\det(W)|$ , where  $W$  is the weight matrix. It is clear that the contribution by a pixel is  $\log |\det(W)|$ , and it has to be decomposed to obtain individual contribution by each channel. As a suitable candidate, singular values of  $W$  correspond to the contribution from each channel dimension, hence their log value is the individual log-det contribution. The individual log-det term for channels are obtained by,

$$|\det(W)| = \prod_i \sigma_i(W) \Leftrightarrow \log |\det(W)| = \sum_i \log(\sigma_i(W)) \quad (5)$$

where  $\sigma_i(W)$  are the singular values of  $W$ . For affine blocks, same method as RealNVP is adopted.

### 3.2.2 Flow models with ODE based Density Estimators

Recent works on flow models such as Behrmann et al. (2018); Grathwohl et al. (2018); Chen et al. (2019) employ variants of ODE based density estimators. The following formulation is applicable to find per-dimensional likelihood contributions for such flow models. In the above works, the flow is modelled as  $F(x)$ , such that  $z = F(x) = (I + g)(x)$ , where  $g(x)$  is the forward propagation function. The log-det of the jacobian is expressed as a power series,

$$\ln |\det J_F(x)| = \text{tr} (\ln (I + J_g(x))) = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{\text{tr}(J_g^k)}{k}$$

where  $\text{tr}$  denotes the trace. Due to computational constraints, the power series is computed up to a finite number of iterations with the  $\text{tr}(J_g^k)$  term stochastically approximated by hutchinson’s trace estimator,  $\text{tr}(A) = \mathbb{E}_{p(v)} [v^T A v]$ , with  $\mathbb{E}[v] = 0$  and  $\text{Cov}(v) = I$ . The component corresponding to each dimension that becomes part of the log-det term is the diagonal element of  $J_g^k$ , so the per-dimension contribution to the likelihood can be approximated as the diagonal elements of  $J_g^k$ , summed over the power series upto a finite number of iterations  $n$ . The diagonal elements are obtained with the hutchinson’s trace estimator without any extra cost, i.e. if  $v = [v_1, v_2, \dots, v_{s \times s \times c}]^T$ ,

$$\sum_{k=1}^{\infty} (-1)^{k+1} \frac{\text{tr}(J_g^k)}{k} = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{\mathbb{E}_{p(v)} [v^T J_g^k v]}{k} = \sum_{k=1}^{\infty} (-1)^{k+1} \frac{\mathbb{E}_{p(v)} [(v^T J_g^k) v]}{k}$$

In the above equation,  $(v^T J_g^k)$  is the vector-jacobian product which is multiplied again with  $v$ . The individual components which are summed when  $(v^T J_g^k)$  is multiplied with  $v$  correspond to the diagonal terms in jacobian, over the expectation  $\mathbb{E}_{p(v)}$ . So those terms are the contribution by the individual dimensions to the log-likelihood and are expressed as  $[L_d^{(l)}]_{s \times s \times c}$  at flow layer  $l$ .

### 3.3 Dimension Factorization using Proposed Heuristic

At every flow layer, an ideal factorization method should,

1. *(Quantitative) For efficient density estimation*: Early gaussianize the dimensions having less log-det and expose the ones having more log-det to more flow layers. In this manner, selectively more expressivity can be provided to dimensions which capture meaningful representations (and are more valuable from a log-det perspective).
2. *(Qualitative) For qualitative reconstruction*: Capture the local variance over the flow layers, i.e. the dimensions being exposed to more flow layers should contain representative pixel variables from throughout the whole image.

We perform a hybrid dimension factorization taking both of the above criterias into account. A  $s \times s \times c$  tensor is converted to  $\frac{s}{2} \times \frac{s}{2} \times 4c$  tensor, by using local max and min pooling operations on corresponding per dimensional log-det terms as obtained in Sec. 3.2 (which was averaged over the training set, so the learned splitting remains same for all data points). Then the  $\frac{s}{2} \times \frac{s}{2} \times 4c$  tensor is split along channel

---

**Algorithm 1:** LCMA Implementation and Training for Generative Flow models

---

**Pre-training Phase:** Pre-train a flow model with no multi-scale architecture (no dimensionality reduction) to obtain the log-det terms ( $[L_d^{(l)}]_{s \times s \times c}$ ) at each layer  $l$

**Dimension Factorization Phase:** Initialize the input dimensions  $\phi \times \phi \times \psi \leftarrow s \times s \times c$

**while**  $1 \leq l \leq K$  **do**

Select  $[L_d^{(l)}]_{\phi \times \phi \times \psi}$  corresponding to input dimensions.

$[L_d^{(l)}]_{\phi \times \phi \times \psi} \xrightarrow[\text{(Figure 1)}]{\text{Local Max- \& Min-Pooling}} [L_d^{(l)}]_{\frac{\phi}{2} \times \frac{\phi}{2} \times 4\psi} \xrightarrow[\text{Splitting}]{\text{Channelwise}} [L_d^{(l, \text{Max})}, L_d^{(l, \text{Min})}]_{\frac{\phi}{2} \times \frac{\phi}{2} \times 2\psi}$

Gaussianize the dimensions corresponding to  $L_d^{(l, \text{Min})}$  and pass the dimensions corresponding to  $L_d^{(l, \text{Max})}$  to more flow layers

$\phi \times \phi \times \psi \leftarrow \frac{\phi}{2} \times \frac{\phi}{2} \times 2\psi$

**end**

**Training Phase:** Flow model with proposed LCMA is trained using maximum likelihood.

---

dimension to form two  $\frac{s}{2} \times \frac{s}{2} \times 2c$  tensors, one corresponding to low log-det dimensions and one corresponding to high ones. The local min and max pooling operations (illustrated in Fig 1) preserve the local spatial variation of the image in both parts of the factorization, leveraging both enhanced density estimation as well as qualitative reconstruction.

**Factorization for Flows involving Squeezing Operation:**

If squeezing operation (which is nothing but reordering of dimensions) is involved, we keep track of which dimensions belong to the half that get gaussianized early and which dimensions belong to the other half that passes through more flow layers. At the next layer, only the log-det terms for the dimensions which came through flow layers are considered for further splitting operation.

We refer to the multi-scale architecture obtained by the above data-dependent dimension splitting method as Likelihood Contribution based Multi-scale Architecture (LCMA).




Figure 1: Likelihood contribution based squeezing operation: (On left) The tensor  $[L_d^{(l)}]_{s \times s \times c}$  representing log-det of variables in a flow layer. (On right) It is squeezed to  $\frac{s}{2} \times \frac{s}{2} \times 4c$  with local max and min pooling operation. The green (orange) marked pixels represent dimensions having more (less) log-det locally.

### 3.4 Integration of LCMA with Flow Model Training

Training of flow models with LCMA implementation is summarized in Algorithm 1.

**Pre-training Phase** involves training a non-multi-scale (NMS) flow model to obtain the individual contribution of dimensions ( $[L_d^{(l)}]_{s \times s \times c}$ ) towards the total log-likelihood (Section 3.2). Training a NMS model is computation heavy, but better training leads to improved density estimation score for resulting LCMA. So, there remains a trade-off between the amount of NMS pre-training and the density estimation performance for resulting LCMA. We train the NMS model partially (we report the results for varying levels of NMS pre-training and resulting density estimation performance for one of the flow training in the supplementary materials).

**Dimension Factorization Phase** involves deciding the dimensions to be factored out at each flow layer using proposed log-det heuristic. The dimension splittings are decided recursively for all the flow layers based on their likelihood contribution obtained in the previous phase.

**Training Phase** is the final stage where the flow model with LCMA is trained. Since the decision for factorization of dimensions at each flow layer occurs before the training starts, and the decision remains unchanged during training, change of variables formula can be applied. In fact, this allows the use of non-invertible operations (e.g. max and min pooling) for efficient factorization (Sec. 3.3).

**Time Complexity:** The time complexity associated with original multi-scale architecture is not affected by LCMA implementation as no additional time consuming blocks were added. Our data-dependent dimension splitting operation can be interpreted as replacing the conventional checkerboard/channel split masking with likelihood contribution based masking.

## 4 Related Work

For multi-scale architectures in generative flow models, our proposed method performs factorization of dimensions based on their likelihood contribution, which in another sense translates to determining which dimensions are important from density estimation and qualitative reconstruction point of view. Keeping this in mind, we discuss prior works on generative flow models which involve multi-scaling and/or incorporate permutation among dimensions to capture their interactions.

A number of generative flow models implement a multi-scale architecture, such as Dinh et al. (2016); Kingma and Dhariwal (2018); Atanov et al. (2019); Izmailov et al. (2019); Durkan et al. (2019); Behrmann et al. (2018); Chen et al. (2019) etc. Kingma and Dhariwal (2018) introduce a  $1 \times 1$  convolution layer in between the actnorm and affine coupling layer in their flow architecture. The  $1 \times 1$  convolution is a generalization of permutation operation which ensures that each dimension can affect every other dimension. This can be interpreted as redistributing the contribution of dimensions to total likelihood among the whole space of dimensions. So Kingma and Dhariwal (2018) treat the dimensions as equiprobable for factorization in their implementation of multi-scale architecture, and split the tensor at each flow layer evenly along the channel dimension. We, on the other hand, take the next step and focus on the individuality of dimensions and their importance from the amount they contribute towards the total log-likelihood. The log-det score is available via direct/indirect decomposition of the jacobian obtained as part of computations in a flow training, so we essentially have an easily available heuristic. Since LCMA focuses individually on the dimensions using easily available heuristic, it can prove to be versatile in being compatible with generic multi-scale architectures. Hoogeboom et al. (2019) extend the concept of  $1 \times 1$  convolutions to invertible  $d \times d$  convolutions, but do not discuss multi-scaling. Dinh et al. (2016) also mention a type of permutation which is equivalent to reversing the ordering of the channels, but is more restrictive and fixed.

Flow models such as Behrmann et al. (2018); Grathwohl et al. (2018); Chen et al. (2019) involve ODE based density estimators. They also implement a multi-scale architecture, but the splitting operation is a static channel wise splitting without considering the importance of individual dimensions or any permutations. Izmailov et al. (2019); Durkan et al. (2019); Kumar et al. (2019); Atanov et al. (2019) use multi-scale architecture in their flow models, coherent with Dinh et al. (2016); Kingma and Dhariwal (2018), but still perform the factorization of dimensions using static masking. For qualitative sampling along with efficient density estimation, we also propose that factorization methods should preserve spatiality of the image in the two splits, motivated by the spatial nature of splitting methods in Kingma and Dhariwal (2018) (channel-wise splitting) and Dinh et al. (2016) (checkerboard and channel-wise splitting). Summarizing, we propose a data-dependent approach to dimension factorization in a multi-scale architecture, unexplored by prior works.

## 5 Experiments

In this section we present the detailed results of proposed LCMA adopted for the flow model of RealNVP (Dinh et al., 2016), Glow (Kingma and Dhariwal, 2018), i-ResNet (Behrmann et al., 2018) and Residual Flows (Chen et al., 2019). For direct comparison, all the experimental settings such as data pre-processing, optimizer parameters as well as flow architectural details (coupling layers, residual blocks) are kept the same, only the factorization of dimensions at each flow layer is performed as per LCMA. The computations were performed in NVIDIA Tesla V100 GPUs. For RealNVP, we perform experiments on four benchmarked image datasets: *CIFAR-10* (Krizhevsky, 2009), *Imagenet* (Russakovsky et al., 2014) (downsampled to  $32 \times 32$  and  $64 \times 64$ ), and *CelebFaces Attributes (CelebA)* (Liu et al., 2015). The scaling in LCMA is performed once for CIFAR-10, thrice for Imagenet  $32 \times 32$  and 4 times for Imagenet  $64 \times 64$  and CelebA. We compare LCMA with conventional RealNVP and report the quantitative and qualitative results. For Glow, i-ResNet and Residual Flows with LCMA, we perform experiments on CIFAR-10 and report improvements over baseline bits/dim (BPD).

### 5.1 Quantitative Comparison

The bits/dim scores of RealNVP with conventional multi-scale architecture and RealNVP with LCMA are given in Table 1. It can be observed that the density estimation results using LCMA is in all cases better in comparison to the baseline. We observed that the

Table 1: Improvements in density estimation (in bits/dim) using proposed method for RealNVP

Model	CelebA	CIFAR-10	ImageNet 32x32	ImageNet 64x64
RealNVP	3.02	3.49	4.28	3.98
RealNVP with LCMA	2.71	3.43	4.21	3.92

improvement for CelebA is relatively high as compared to natural image datasets. This observation was expected as facial features often contain high redundancy and the flow model learns to put more importance (reflected in terms of high log-det) on selected dimensions that define the facial features. Proposed LCMA exposes such dimensions to more flow layers, making them more expressive and hence the significant improvement in BPD is observed. The improvement in bits/dim is less for natural image datasets because of the high variance among features defining them, which has been the challenge with image compression algorithms. Note that the improvement in density estimation is always relative to the original flow architecture (RealNVP in our case) over which we use our proposed LCMA, as we do not alter any architecture other than the dimension factorization method.

The quantitative results of LCMA implementation for several state-of-the-art flow models with CIFAR-10 dataset is summarized in Table 2. The density estimation score for flow with LCMA outperforms the same flow with conventional multi-scale architecture. We also achieve state-of-the-art density estimation results for CIFAR-10 dataset with LCMA implementation for Residual Flows.

Table 2: Density estimation results (in bits/dim) for various flow models with LCMA on CIFAR-10

Type of Multi-scale Architecture (MA)	RealNVP	Glow	i-ResNet	Residual Flows
Conventional MA	3.49	3.35	3.45	3.28
LCMA	<b>3.43</b>	<b>3.31</b>	<b>3.40</b>	<b>3.25</b>

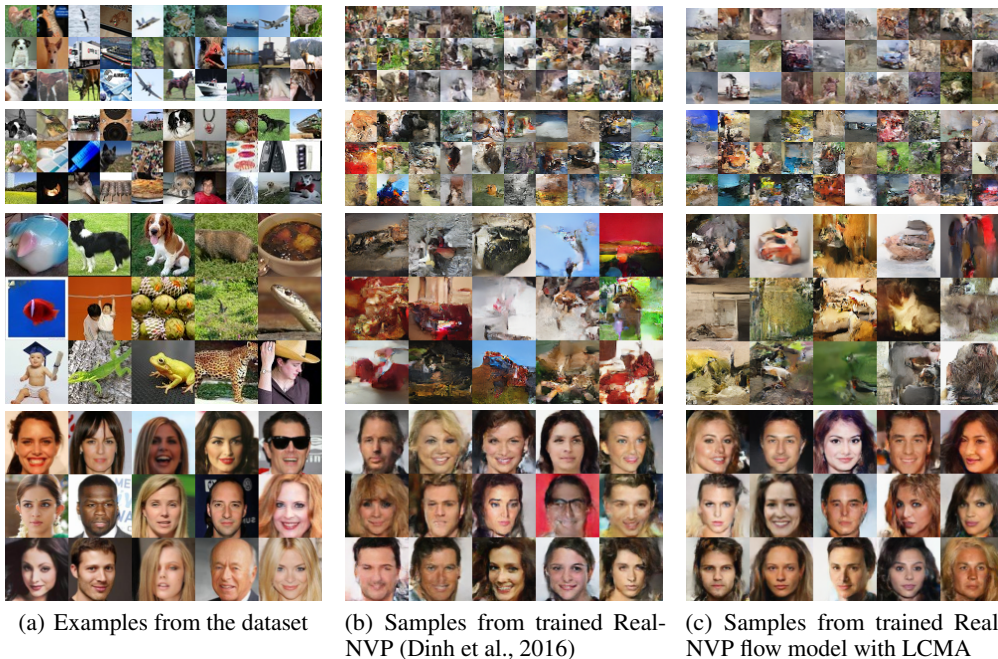


Figure 2: Samples from RealNVP (Dinh et al., 2016) and RealNVP flow model with proposed LCMA trained on different datasets. The datasets shown in this figure are in order: CIFAR-10, Imagenet(32 × 32), Imagenet (64 × 64) and CelebA (without low-temperature sampling).

## 5.2 Qualitative Comparison

For LCMA implementation, we introduced local max and min pooling operations (to preserve spatiality) on log-det heuristic to decide which dimensions to be gaussianized early (Section 3). Fig. 2(a) shows samples from original datasets, Fig. 2(b) shows the samples from a trained RealNVP flow model with conventional multi-scale architecture and Fig. 2(c) shows the samples from RealNVP with LCMA, trained on various datasets. The finer facial details such as hair styles, eye-lining and facial folds in Celeba samples generated from RealNVP with LCMA were perceptually better than the baseline. The global



Figure 3: Smooth linear interpolations in latent space between two images from CelebA. The intermediate samples perceptibly resemble synthetic faces.







feature representation observed is similar to that in RealNVP, as the flow architecture was kept the same. The background for natural images such as Imagenet were constructed at par with the original flow model. As it has been observed, for flow models, the latent space holds knowledge about the feature representation in the data. We performed linear interpolations in latent space to ensure its efficient construction. The interpolations observed (Fig. 3) were smooth, with intermediate samples perceptibly resembling synthetic faces, signifying the efficient construction of latent space.

### 5.3 Ablation Study

We perform two types of ablation studies to compare LCMA with other methods for dimension factorization in a multi-scale architecture. In our first study, we consider 4 variants, namely fixed random permutation (Case 1), multiscale architecture with early gaussianization of high log-det dimensions (Case 2), factor-

Table 3: Ablation study results for multi-scale architectures with various factorization methods trained on CelebA dataset

Evaluations	Fixed Random Permutation	Early gaussianization of high log-det dimensions	RealNVP	Early gaussianization of low log-det dimensions
Quantitative (BPD)	3.05	3.10	3.02	2.71
Qualitative				

ization method with checker-board and channel splitting as introduced in RealNVP (Case 3) and multiscale architecture with early gaussianization of low log-det dimensions, which is our proposed LCMA (Case 4). In fixed random permutation, we randomly partition the tensor into two halves, with no regard to the spatiality or log-det score. In case 2, we do the reverse of LCMA, and early gaussianize the high log-det variables at each layer. The bits/dim score and generated samples for each of the methods are given in Table 3. As expected from an information theoretic perspective, gaussianizing high log-det variables early provides the worst density estimation, as the model could not capture the high amount of important information. Comparing the same with fixed random permutation, the latter has better score as the probability of a high log-det variable being gaussianized early reduces to half, and it gets further reduced with RealNVP due to channel-wise and checkerboard splitting. LCMA has the best score among all methods, as the variables more sensitive to changes in input (hence carrying more information) are exposed to more flow layers. Fixed random permutation has the worst quality of sampled images, as the spatiality is lost during factorization. The sample quality improves for Case 2 and RealNVP. The sampled images are perceptually best for LCMA.

We perform a second ablation study to reconfirm that early gaussianization of high log-det dimensions has a deteriorating effect on the density estimation score. The flow model in our experiment has 3 layers where dimensions splitting is being performed. We consider all permutations of early gaussianization of high/low log-det dimensions at each of the 3 layers. The density estimation scores for all  $2^3$  permutations trained on CelebA dataset are presented in Table 4. The best score corresponds to early gaussianization of low log-det dimensions at each flow layer (proposed LCMA), and the score deteriorates with permutations involving early gaussianization of high log-det dimensions at any flow layer. Summarizing, LCMA outperforms multi-scale architectures based on other factorization methods, as it improves density estimation scores and generates qualitative samples.

Table 4: Ablation study results for permuting factorization of high/low log-det dims

Permutation of high/low log-det dimensions	Bits/dim
High-High-High	3.10
High-High-Low	3.09
High-Low-High	3.07
High-Low-Low	3.05
Low-High-High	3.00
Low-High-Low	2.92
Low-Low-High	2.79
Low-Low-Low (LCMA)	2.71

## 6 Conclusions

We proposed a novel multi-scale architecture for generative flows which employs a data-dependent splitting based on the individual contribution of dimensions to the log-likelihood. Implementations of the proposed method for several flow models such as RealNVP (Dinh et al., 2016), Glow(Kingma and Dhariwal, 2018), i-ResNet(Behrmann et al., 2018) and Residual Flows (Chen et al., 2019) were presented. Empirical studies conducted on benchmark image datasets validate the strength of our proposed method, which improves log-likelihood scores and is able to generate qualitative samples. Ablation study results confirm the power of LCMA over other options for dimension factorization. A line of future work can be to design/learn a masking scheme for factorization online during flow training (or possibly a parallel training process), while preserving flow properties.

## Broader Impact

Generative models in general require higher amounts of computations in order to be trained with a dense network and/or larger amounts of data. This becomes a big challenge for real-world high-dimensional data, especially when the model itself has constraints of bijectivity, e.g. generative flows. Proposed Likelihood Contribution based Multi-scale Architecture (LCMA) provides an avenue to evaluate the importance of individual dimensions present in the data and use it to make flow based generative models computationally efficient and at the same time powerful density estimators.

We envision proposed research will open up discussions on making flow based generative models computationally efficient, so as to allow a broad section of disciplines reap benefits of it for their big-data applications while remaining within their computational budgets. Particularly this will make researchers with limited budget get access and contribute to the generative AI world. Optimized flow models can also be used alongside real-world applications such as energy/smart buildings (Zou et al., 2019c,b,a; Konstantakopoulos et al., 2019; Chen et al., 2021; Periyakoil et al., 2021, 2020; Das et al., 2019, 2020; Liu, 2018; Liu et al., 2019; Donti and Kolter, 2021; Jin et al., 2018; Das et al., 2021a), healthcare (Das et al., 2021c), vision (Das et al., 2021b; Yue et al., 2021; Gong et al., 2019) etc. Additionally, since we focus on quantifying the sensitivity of each data dimension to the input data variations, we expect our method will potentially be used in data compression, semi-supervised learning and to deal with data bias, which are important components in dealing with AI for societal systems.

There still are directions where the proposed research has room to be more efficient (which we also mention as future work). Currently, proposed dimension splitting occurs offline (not during the actual training process), since doing it otherwise would violate *change of variables* formula, a general principle for flow training. A design of dimension splitting which directly/indirectly occurs online (during the training process) and is in line with flow training principles can make the method even more computationally efficient, hence more accessible to researchers with limited budget.

## Acknowledgement

This research was funded by the Republic of Singapore’s National Research Foundation through a grant to the Berkeley Education Alliance for Research in Singapore (BEARS) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) Program. BEARS has been established by the University of California, Berkeley as a center for intellectual excellence in research and education in Singapore.

## References

- A. Atanov, A. Volokhova, A. Ashukha, I. Sosnovik, and D. Vetrov. Semi-conditional normalizing flows for semi-supervised learning. *arXiv preprint arXiv:1905.00505*, 2019.
- J. Behrmann, W. Grathwohl, R. T. Chen, D. Duvenaud, and J.-H. Jacobsen. Invertible residual networks. *arXiv preprint arXiv:1811.00995*, 2018.
- B. Chen, P. Donti, K. Baker, J. Z. Kolter, and M. Berges. Enforcing policy feasibility constraints through differentiable projection for energy optimization. *arXiv preprint arXiv:2105.08881*, 2021.
- R. T. Chen, J. Behrmann, D. Duvenaud, and J.-H. Jacobsen. Residual flows for invertible generative modeling. *arXiv preprint arXiv:1906.02735*, 2019.
- X. Chen, N. Mishra, M. Rohaninejad, and P. Abbeel. Pixelsnail: An improved autoregressive generative model. *arXiv preprint arXiv:1712.09763*, 2017.
- M. Cisse, P. Bojanowski, E. Grave, Y. Dauphin, and N. Usunier. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 854–863. JMLR. org, 2017.
- H. P. Das, I. C. Konstantakopoulos, A. B. Manasawala, T. Veeravalli, H. Liu, and C. J. Spanos. A novel graphical lasso based approach towards segmentation analysis in energy game-theoretic frameworks. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, pages 1702–1709. IEEE, 2019.

- H. P. Das, I. Konstantakopoulos, A. B. Manasawala, T. Veeravalli, H. Liu, and C. J. Spanos. Do occupants in a building exhibit patterns in energy consumption? analyzing clusters in energy social games. *Workshop on Tackling Climate Change with Machine Learning, Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- H. P. Das, S. Schiavon, and C. J. Spanos. Unsupervised personal thermal comfort prediction via adversarial domain adaptation. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pages 230–231, 2021a.
- H. P. Das, R. Tran, J. Singh, Y.-W. Lin, and C. J. Spanos. Cdcgen: Cross-domain conditional generation via normalizing flows and adversarial training, 2021b.
- H. P. Das, R. Tran, J. Singh, X. Yue, G. Tison, A. Sangiovanni-Vincentelli, and C. J. Spanos. Conditional synthetic data generation for robust machine learning applications with limited pandemic data. In *arXiv preprint arXiv:2109.06486*, 2021c.
- L. Dinh, D. Krueger, and Y. Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real NVP. *CoRR*, abs/1605.08803, 2016.
- P. L. Donti and J. Z. Kolter. Machine learning for sustainable energy systems. *Annual Review of Environment and Resources*, 46, 2021.
- C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios. Cubic-spline flows. *arXiv preprint arXiv:1906.02145*, 2019.
- J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts. Gansynth: Adversarial neural audio synthesis. *arXiv preprint arXiv:1902.08710*, 2019.
- R. Gong, W. Li, Y. Chen, and L. V. Gool. Dlow: Domain flow for adaptation and generalization, 2019.
- W. Grathwohl, R. T. Chen, J. Betterncourt, I. Sutskever, and D. Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.
- J. Ho, X. Chen, A. Srinivas, Y. Duan, and P. Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. *arXiv preprint arXiv:1902.00275*, 2019.
- E. Hooeboom, R. v. d. Berg, and M. Welling. Emerging convolutions for generative normalizing flows. *arXiv preprint arXiv:1901.11137*, 2019.
- P. Izmailov, P. Kirichenko, M. Finzi, and A. G. Wilson. Semi-supervised learning with normalizing flows. In *Workshop on Invertible Neural Nets and Normalizing Flows, International Conference on Machine Learning*, 2019.
- M. Jin, R. Jia, H. P. Das, W. Feng, and C. Spanos. Biscuit: Building intelligent system customer investment tool. In *10th International Conference on Applied Energy (ICAE)*, 2018.
- T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*, 2018.
- D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pages 3581–3589, 2014.

- I. C. Konstantakopoulos, H. P. Das, A. R. Barkan, S. He, T. Veeravalli, H. Liu, A. B. Manasawala, Y.-W. Lin, and C. J. Spanos. Design, benchmarking and explainability analysis of a game-theoretic framework towards energy efficiency in smart infrastructure. *arXiv preprint arXiv:1910.07899*, 2019.
- A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- M. Kumar, M. Babaeizadeh, D. Erhan, C. Finn, S. Levine, L. Dinh, and D. Kingma. Videoflow: A flow-based generative model for video. *arXiv preprint arXiv:1903.01434*, 2019.
- S. Liu. Personal thermal comfort models based on physiological parameters measured by wearable sensors. 2018.
- S. Liu, S. Schiavon, H. P. Das, M. Jin, and C. J. Spanos. Personal thermal comfort models with wearable sensors. *Building and Environment*, 162:106281, 2019.
- Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015.
- A. Odena. Semi-supervised learning with generative adversarial networks. *arXiv preprint arXiv:1606.01583*, 2016.
- A. V. d. Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016a.
- A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016b.
- D. Periyakoil, H. P. Das, and C. J. Spanos. Understanding distributions of environmental parameters for thermal comfort study in singapore. In *Proceedings of the Eleventh ACM International Conference on Future Energy Systems*, pages 461–465, 2020.
- D. Periyakoil, H. P. Das, C. Miller, C. J. Spanos, and N. Prata. Environmental exposures in singapore schools: An ecological study. *International journal of environmental research and public health*, 18(4):1843, 2021.
- O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li. Imagenet large scale visual recognition challenge. *CoRR*, abs/1409.0575, 2014. URL <http://arxiv.org/abs/1409.0575>.
- T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *CoRR*, abs/1701.05517, 2017. URL <http://arxiv.org/abs/1701.05517>.
- B. Uria, I. Murray, and H. Larochelle. Rnade: The real-valued neural autoregressive density-estimator. In *Advances in Neural Information Processing Systems*, pages 2175–2183, 2013.
- X. Yue, Z. Zheng, H. P. Das, K. Keutzer, and A. S. Vincentelli. Multi-source few-shot domain adaptation. *arXiv preprint arXiv:2109.12391*, 2021.
- H. Zou, H. P. Das, J. Yang, Y. Zhou, and C. Spanos. Machine learning empowered occupancy sensing for smart buildings. *Climate Change + AI Workshop, International Conference on Machine Learning (ICML)*, 2019a.
- H. Zou, J. Yang, H. Prasanna Das, H. Liu, Y. Zhou, and C. J. Spanos. Wifi and vision multimodal learning for accurate and robust device-free human activity recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019b.
- H. Zou, Y. Zhou, J. Yang, H. Liu, H. P. Das, and C. J. Spanos. Consensus adversarial domain adaptation. In *AAAI Conference on Artificial Intelligence 2019*, 2019c.