

A Forest from the Trees: Generation through Neighborhoods

Yang Li, Tianxiang Gao, Junier B. Oliva

Department of Computer Science, UNC Chapel Hill

{yangli95, tianxiang, joliva}@cs.unc.edu

Abstract

In this work, we propose to learn a generative model using both learned features (through a latent space) and memories (through neighbors). Although human learning makes seamless use of both learned perceptual features and instance recall, current generative learning paradigms only make use of one of these two components. Take, for instance, flow models, which learn a latent space that follows a simple distribution. Conversely, kernel density techniques use instances to shift a simple distribution into an aggregate mixture model. Here we propose multiple methods to enhance the latent space of a flow model with neighborhood information. Not only does our proposed framework represent a more human-like approach by leveraging both learned features and memories, but it may also be viewed as a step forward in non-parametric methods. In addition, our proposed framework allows the user to easily control the properties of generated samples by targeting samples based on neighbors. The efficacy of our model is shown empirically with standard image datasets. We observe compelling results and a significant improvement over baselines. Combined further with a contrastive training mechanism, our proposed methods can effectively perform non-parametric novelty detection.

1 Introduction

The typical training paradigm for many generative models makes a one time use of the training data. For instance, VAEs (Kingma and Welling 2013) and flow models (Dinh, Krueger, and Bengio 2014; Dinh, Sohl-Dickstein, and Bengio 2016; Kingma and Dhariwal 2018) use the training data to fit a latent space. After the features (the latent space) are learned, the training data is discarded, as the learned network is entirely responsible for the generative process. This paradigm has proven effective, however it leaves the entire burden of modeling a complicated support (such as the space of images) completely on the learned latent space and network capacity. Moreover, such an approach is in stark contrast to human learning, which not only uses data to learn perceptual features (Kuhl, Tsao, and Liu 2003), but will reuse data as memories with instance recall (Carrier and Pashler 1992). Furthermore, since samples are generated from a random latent code, it

is difficult to control the properties of generated samples. In this work, we more closely follow nature’s paradigm for generative modeling, making use of our data not only to learn features but to pull instances as neighbors.

Neighborhood based methods abound much of the history of machine learning and statistics. For instance, the well-known k -nearest neighbor estimator often achieves surprisingly good results (Altman 1992). Moreover, by equipping models with neighborhood information, one reduces the problem to a local-manifold, a much simpler domain (Roweis and Saul 2000; Seung and Lee 2000). In this work, we focus our attention on density estimation, which features a staple neighbor based estimator: the kernel density estimator (KDE). The KDE models a local manifold as a Gaussian, an untenable restriction when dealing with complicated data such as images, see Fig. 1. Instead, we model local manifolds using more flexible densities, simultaneously better representing the local space of a neighborhood whilst reducing the burden of the density estimator. In doing so, this work brings forth methodology that is in the same vein as non-parametric density estimation. We begin with a simple observation about the kernel density estimator. Recall that KDE models a dataset $\{x_i \in \mathbb{R}^d\}_{i=1}^N$ as a mixture of “kernel” distributions K centered around each data-point:

$$p(x) = \frac{1}{N} \sum_{i=1}^N K(x_i, x). \quad (1)$$

In practice, the kernel is often chosen to be a Gaussian distribution centered at the data-points with a diagonal covariance matrix:

$$p(x) = \frac{1}{N} \sum_{i=1}^N \mathcal{N}(x | x_i, \sigma^2 \mathbf{I}), \quad (2)$$

where $\mathcal{N}(x | x_i, \sigma^2 \mathbf{I})$ is the pdf of a Gaussian with mean x_i and covariance matrix $\sigma^2 \mathbf{I}$. That is, the density is modeled as a mixture of training data-centered distributions. The sampling procedure is simple, as one selects a training point x_i from the training data at random and adds Gaussian noise to it. More generally, one sees that KDE models the density as:

$$p(x) = \frac{1}{N} \sum_{i=1}^N p(x | x_i), \quad (3)$$

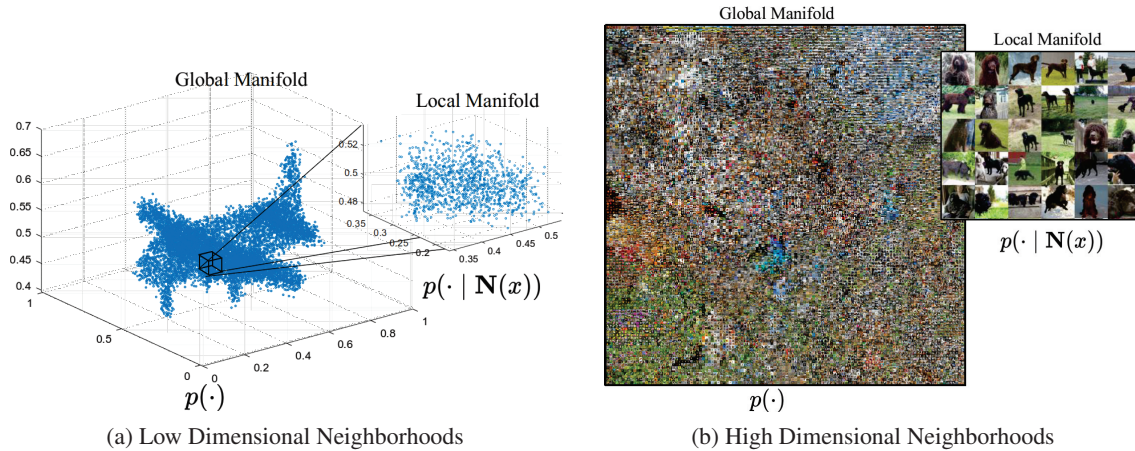


Figure 1: Neighborhood based models, which focus on modeling a simpler local manifold rather than the more complicated global support. (a) In low dimensions, one may model tight local manifolds using simple distributions (e.g. Gaussians). Thus, in low dimensions it is relatively simple to model a rich class of densities as the mixture of shifted simple distributions (i.e. kernels). (b) In higher dimensions, such as the space of natural images, the number of simple kernels to adequately model the space grows exponentially, and is thus intractable. Tractable local manifolds may no longer be Gaussian, however they still exhibit a structure that can be exploited. In this work, we propose to model local manifolds using flexible density estimators and careful conditioning.

for a very restricted class of conditionals $p(\cdot|x_i)$, namely shifted Gaussians. We propose a data-driven generalization of KDE, another step forward in non-parametric methods. We use neighborhoods to allow the estimator to model local manifolds, a simpler space. Specifically, we model the density as:

$$p(x) = \frac{1}{N} \sum_{i=1}^N p_{\theta}(x | \mathbf{N}(x_i)), \quad (4)$$

where $p_{\theta}(\cdot | \mathbf{N}(x_i))$ is a more flexible density estimator conditioned on neighborhood information (Fig. 1(b)) and $\mathbf{N}(x_i)$ is the neighborhood around point x_i .

Our approach (4) is capable of inferring relevant features and variances of the neighborhood, lessening the burden on p_{θ} . Note further that KDE (2) blurs the line between density estimation and data augmentation; KDE samples by “augmenting” a data point with Gaussian noise. The proposed approach expands on this by learning how to “augment” or sample given a neighbor or neighborhood.

The sampling procedure for our proposed model would remain simple: one would choose a training point x_i uniformly at random from the training data, then sample according to the estimator conditioned on the neighborhood: $x \sim p_{\theta}(\cdot | \mathbf{N}(x_i))$. Note that the entire sampling procedure is not a conditional model since we marginalize out the effect of the neighborhood by picking them from training data at random (4).

In addition to the simple global unconditioned sampling procedure, our paradigm also offers an unprecedented targeted sampling interface. Suppose that a user wanted to produce novel images of a tiger in a grassy outdoor setting, standing, and facing the camera. In order to provide a user with such fine-grained control in standard conditional models,

one would have to train a multi-labeled model based on concepts such as “contains tiger”, “orient forward”, “outdoor”, “grassy”, etc. However, such a labeled dataset may be very expensive to collect. Furthermore, it may be hard to obtain a compact representation for certain concepts (Higgins et al. 2016). Lastly, a multi-labeled conditional model would be unable to extrapolate to unforeseen concepts at test times. These shortcomings make a traditional approach unfit for a targeted sampling task. Instead, our neighbor based approach is very adept at producing targeted novel samples. In the aforementioned scenario, the user would be able to condition the neighborhood model using tiger images similar to his/her target. In effect, our model provides an indispensable “more like this” interface to generative models.

Our proposed approach builds on the standard non-parametric methodology in three key ways. *First*, we propose to use more robust density kernels than the standard approaches; while kernels like the Gaussian distribution lead to general distributions asymptotically, their simple unimodal nature often falls short for higher dimensions in finite samples (Wasserman 2010). *Second*, we propose to condition on training data in a more robust fashion than by simply shifting a base distribution as is standard; this allows one to capture a richer set of correlations of data in a local neighborhood. *Third*, we propose to extract the set-level information from the neighborhood $\mathbf{N}(x_i)$ in a data-driven fashion; conditioning on multiple instances in a neighborhood allows our model to learn the variances present in the local manifold (in the background of images, for instance) without the network needing to memorize them.

In all, our proposed paradigm makes use of both memories (via instance recall of neighborhoods) and learned features (via a latent space). As such, our approach is richer than

existing methods, more closely resembles human learning, and gives us fine-grained control over sample properties. We show the efficacy of our models on standard benchmark image datasets with experiments below.

2 Methods

While the local manifold of a neighborhood might be remarkably simpler than the global manifold, one is still unlikely to estimate the local manifold well with a simple base distribution such as a Gaussian. Thus, we still need to capture the data in these local neighborhoods \mathbf{N} with a flexible model $p_\theta(\cdot | \mathbf{N})$. The choice of robust base conditional model is quite flexible, with many possible choices including autoregressive models (Oord, Kalchbrenner, and Kavukcuoglu 2016; van den Oord et al. 2016), flow models (Dinh, Krueger, and Bengio 2014; Dinh, Sohl-Dickstein, and Bengio 2016; Kingma and Dhariwal 2018), and variational autoencoders (VAEs) (Kingma and Welling 2013). In fact, one may also deviate from a likelihood based framework in favor of a discriminative critic GAN approach (Goodfellow et al. 2014). As our focus is on the methodology of injecting neighborhood information into a generative approach, we explore this with a particular choice of generative model, flow density estimators. However, our approach easily extends to other generative models. In addition, flow methods provide a few advantages including: a tractable (normalized) likelihood, robustness to mode collapse (Grover, Dhar, and Ermon 2018), and a meaningful latent space. We expound on the base flow generative models below.

2.1 Flow based Models

The change of variable theorem, shown in (5), is the cornerstone of flow generative models (Dinh, Krueger, and Bengio 2014; Dinh, Sohl-Dickstein, and Bengio 2016; Kingma and Dhariwal 2018), where q is an invertible transformation.

$$p_X(x) = \left| \det \frac{dq}{dx} \right| p_Z(q(x)) \quad (5)$$

In order to efficiently compute the determinant, the transformation q is often designed to have diagonal or triangular Jacobian. Since this type of transformation is restricted, the flow models often compose multiple transformations in a sequence to get a flexible transformation, i.e. $q = q_m \circ q_{m-1} \circ \dots \circ q_1$. Here, the covariates flow through a chain of transformations, substituting the last output variable as input for the next transformation.

One family of this type of transformation is the so-called *coupling* layer used in NICE (Dinh, Krueger, and Bengio 2014), RealNVP (Dinh, Sohl-Dickstein, and Bengio 2016) and Glow (Kingma and Dhariwal 2018). The input x is divided into two parts x^A and x^B , the first part is kept the same, and the second part is transformed based on the first part using an affine transformation,

$$\begin{aligned} y^A &= x^A \\ y^B &= x^B \odot s(x^A) + t(x^A), \end{aligned} \quad (6)$$

where \odot represents the element-wise (Hadamard) product. Such coupling layers are easy to invert and the Jacobian determinant is easy to compute. Furthermore, since the Jacobian determinant does not involve computing the Jacobian of $s(\cdot)$ and $t(\cdot)$, they can potentially be any deterministic function, such as a neural network.

Since one single coupling layer can only affect part of the input covariates, earlier works swap A and B between two coupling layers to ensure every dimension can affect every other dimension. Recent work (Oliva et al. 2018; Kingma and Dhariwal 2018) proposes to learn a *linear* layer to better capture the correlations along dimensions. In order to stabilize the training and ensure the whole process’s invertibility, the RealNVP reformulates the batch normalization into an invertible transformation. They also propose a multi-scale architecture to capture spatial correlations in image data.

In the end, the flow models transform inputs to a new space, $z = q(x)$, where covariates can be modeled using a simple base distribution, see Fig. 2(a). Typically, we take $p_Z(z)$ to be a Gaussian distribution. From another perspective, we use a deep neural network to construct a series of flexible and invertible operations to transform a simple distribution (e.g. a Gaussian) into a complicated one (the data distribution). Flow based generative models use the exact log-likelihood $\log p(x)$ as the training criterion.

2.2 Neighbor Conditioned Flow Models

We impose the neighborhood information in neighbor conditioned distributions (4) by conditioning a flow model. Flow models can be divided into two parts, the transformation of variables and the latent distribution. Both of these parts can incorporate extraneous conditioning information. Concretely, the transformation procedure can use neighborhood information to decide how to transform the inputs, while the latent distribution can have mean and variance depending on conditioning variables. We explore two ways to inject neighborhood information into the flow model. First, we propose neighborhood conditioned likelihood (NCL), which specifies the distribution of the latent covariates of a target point, $z = q(x)$, given its neighborhood \mathbf{N} . Second, we propose neighborhood conditioned transformations (NCT), which adjusts the latent space (transformation of variables) for a target point x according to its neighborhood \mathbf{N} , i.e., $z = q_{\mathbf{N}}(x)$.

Neighborhood Conditioned Likelihood (NCL) We propose to directly estimate the density of the representation of a target point x in the latent space, $q(x)$, given its neighborhood \mathbf{N} . To do so, two mappings, g_μ , and g_σ , are learned to estimate mean and variance parameters, $\bar{\mu}$ and $\bar{\sigma}$, respectively from \mathbf{N} . We disentangle the transformation of variables from the relationship between the neighbors and the target by passing in transformed neighbors to g_μ, g_σ , i.e., we pass neighbor images through the same flow transformations q . Then, we use g_μ and g_σ to compute Gaussian means and variances respectively. As in original flow models, we use a diagonal matrix for the variance:

$$p_Z(z) = \mathcal{N}(z | g_\mu(q(\mathbf{N})), \text{diag}(g_\sigma^2(q(\mathbf{N})))) \quad (7)$$

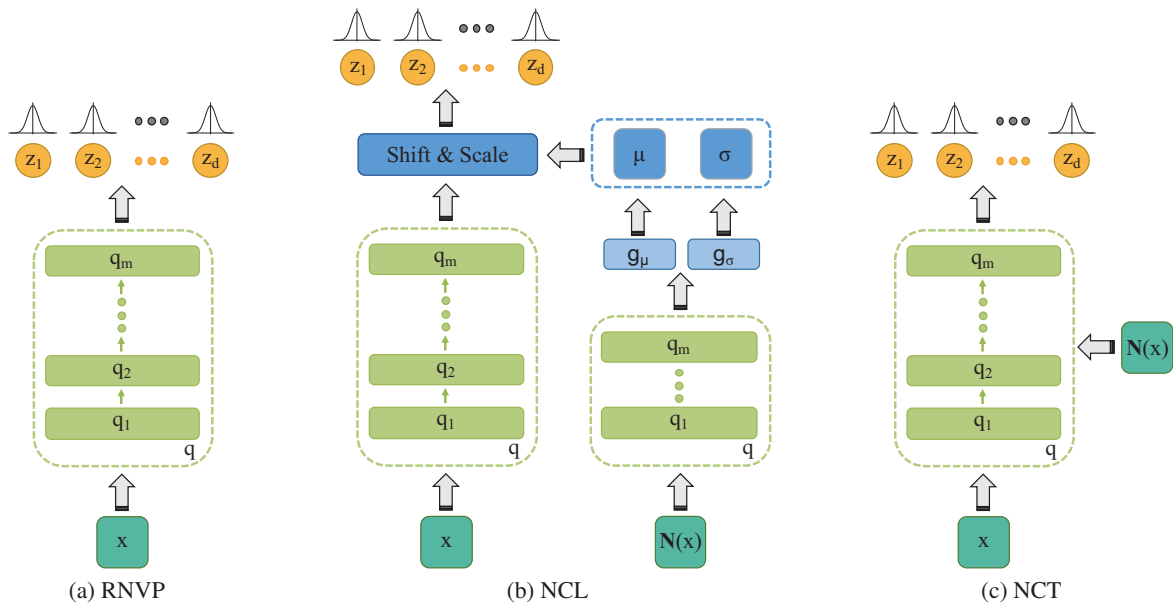


Figure 2: (a) Vanilla RealNVP model. The dashed box q indicates a series of invertible transformations, which transform images $x \in \mathbb{R}^d$ (cyan blocks) into d standard Gaussian covariates (yellow circles). (b) Conditioning latent distribution on neighbors. Note that the two dashed green boxes share the same group of parameters (same transformation for input images and corresponding neighborhoods). g_μ and g_σ compute mean and variance based on the specified neighborhood respectively. (c) Using neighbors to condition the transformations. Some or all transformations in q could depend on neighborhood $\mathbf{N}(x)$.

The model is illustrated in Fig. 2(b). To sample, we choose a neighborhood \mathbf{N} uniformly at random from the training data. Then, we transform each neighbor image in \mathbf{N} into the latent space, and get the neighbor conditioned distribution based on latent codes of neighbors. We then sample latent codes, $z \sim \mathcal{N}(\cdot | g_\mu(q(\mathbf{N})), \text{diag}(g_\sigma^2(q(\mathbf{N}))))$, from this distribution and invert the flow transformation to get image samples, $x = q^{-1}(z)$. In another point of view, our **NCL** model resembles a KDE model in the latent space since we are sampling from a Gaussian conditioned on a neighborhood. However, as we are operating in a latent space and are conditioning on the neighborhood according to the output of a general mapping, the **NCL** is a strictly more general model.

Neighborhood Conditioned Transformations (NCT)

As mentioned, another component of flow models that is amenable to conditioning information is the transformation of variables. Although the **NCL** model allows the latent covariates z to come from a distribution that depends on a neighborhood \mathbf{N} , the construction of z itself is uninformed by \mathbf{N} . We propose to inject neighborhood information into flow transformations using neighborhood conditioned transformations (**NCT**). The **NCT** model provides guidance about how to effectively transform a local manifold through neighbors. Here, we propose a neighborhood conditioned coupling transformation:

$$\begin{aligned} y^A &= x^A \\ y^B &= x^B \odot s(x^A, \mathbf{N}) + t(x^A, \mathbf{N}). \end{aligned} \quad (8)$$

The shift and scale functions s and t are implemented via a concatenation operation on inputs. Replacing the coupling layer in a standard flow model with the proposed neighbor conditioned one gives rise to our **NCT** model. The transformations (8) can model both intra-pixel and intra-neighbor dependencies. Furthermore, by conditioning the flow transformations throughout each of the composing coupling transformation, we are able to inject our conditioning neighborhood information through the transformation of variables, rather than only at the end. In cases of multi-scale architecture, like the one used in RealNVP, x^A could have different spatial dimensions from \mathbf{N} , thus we re-size the neighbors before concatenating them.

We note that the **NCL** model can be interpreted as a special case of **NCT**. One may view g_μ and g_σ as specifying a shift and scale operation in the transformed space:

$$z = \frac{q(x) - g_\mu(q(\mathbf{N}))}{g_\sigma(q(\mathbf{N}))}, \quad (9)$$

where z now can be modeled as isotropic unit norm Gaussian. However, as we specify g_μ and g_σ in terms of the transformation q , we get distinct models stemming from each approach. The details about our neighbor conditioned flow models are presented in the appendix ¹.

2.3 Contrastive Training

A straight-forward way to train our neighbor conditioned models is to directly optimize the neighbor conditioned likeli-

¹please refer to <https://arxiv.org/abs/1902.01435> for a full version of this paper.

hood. However, likelihood trained models are known to suffer from out-of-distribution (OoD) issues (Nalisnick et al. 2018; Choi and Jang 2018), which means such models may assign a higher likelihood to data from a completely different distribution than data it is trained on. When deploying the models to perform novelty/anomaly detection, the insensitivity to OoD data can completely devastate the prediction. Inspired by the noise contrastive estimation (NCE) (Gutmann and Hyvärinen 2010), we propose a contrastive training mechanism to increase the sensitivity of our models to OoD data. A common and major challenge with NCE is the choice of the noise distribution and how to obtain “negative” samples. Fortunately, our neighborhood based paradigm provides a natural answer. For a given neighborhood \mathbf{N} , $p(\cdot | \mathbf{N})$ is to have support only in the local manifold around \mathbf{N} . Hence, we can take $p(\cdot | \neg\mathbf{N})$ for a different neighborhood to be a noise distribution w.r.t. $p(\cdot | \mathbf{N})$ and take $x \in \neg\mathbf{N}$ as negative samples.

That is, given x_i and x_j from two different neighborhoods $\mathbf{N}(x_i)$ and $\mathbf{N}(x_j)$ respectively, we optimize the following objective function:

$$\mathcal{L}_\theta = \mathbb{E}_{x \sim p(\cdot | \mathbf{N}(x_i))} \mathbb{E}_{x' \sim p(\cdot | \mathbf{N}(x_j))} [p_\theta(x | \mathbf{N}(x_i)) + \max(0, m + p_\theta(x' | \mathbf{N}(x_i)) - p_\theta(x | \mathbf{N}(x_i)))] \quad (10)$$

where m represents a margin between in-neighborhood likelihood and out-of-neighborhood likelihood. Although the original NCE objective uses an infinite margin, we find this finite margin objective stabilizes the training.

3 Experiments

Following the exact preprocessing procedure in RealNVP, we transform the pixels into logit space to alleviate the impact of boundary effects. In order to conduct a fair comparison with RealNVP, we use exactly the same network architecture and hyperparameters as those in (Dinh, Sohl-Dickstein, and Bengio 2016). We do not apply data augmentation for all experiments, but we use early stopping to prevent overfitting. We optimize the likelihood objective with contrastive training in novelty detection experiments.

In order to avoid cheating, we pull neighbors only from training set. Please refer to appendix for more information about getting neighborhoods. To explore the sensitivity of our methods to the choice of neighborhoods, we also inspect the behavior of our models using far fewer neighborhoods, where we cluster the training data into several neighborhoods, see appendix for more details.

In addition to vanilla RealNVP, we also compare to a simple class label/attributes conditioned RealNVP model. We model the latent space as a class/attributes conditioned Gaussian distribution. We use one fully connected layer to derive the Gaussian mean and variance respectively. In contrast to the Glow model (Kingma and Dhariwal 2018), we do not share the class conditioned Gaussian over spatial dimensions, which we observed made our class conditioned model a much stronger baseline.

When training based on the likelihood for a specific data point x_i , we optimize log likelihood values conditioned on the respective neighborhood, i.e., $\log p_\theta(x_i | \mathbf{N}(x_i))$. While

Table 1: FID scores. GAN results are from (Lucic et al. 2018; Heusel et al. 2017). Best scores are marked in **bold**.

Dataset	GAN	RNVP	CC-RNVP	NCL	NCT
MNIST	6.7	10.4	8.3	11.2	8.2
SVHN	12.5	103.8	55.2	50.0	61.8
CIFAR-10	55.2	99.9	93.4	95.8	81.1
CelebA	30.0	39.2	32.5	33.0	30.9

the likelihood of the generative process is (4), optimizing our neighborhood criteria encourages likelihoods to be concentrated around the respective conditioning neighborhood. FID scores (Heusel et al. 2017) and precision-recall (PRD) (Sajjadi et al. 2018) are used to quantitatively compare the sample quality. As in (Sajjadi et al. 2018), we report $(F_8, F_{\frac{1}{8}})$ pair to represent the recall and precision respectively. Models are compared using 50K samples.

3.1 Results and Analysis

We show samples generated for all four datasets in Fig. 3. We also show actual images in the first column for comparison.

On all datasets we see that our models produce coherent samples with respect to the conditioned neighbors. For example, on MNIST our neighbor based models allow fine-grained control over properties of samples, such as writing style and the brush width. In SVHN, the vanilla RealNVP yields lesser quality samples, and the class conditioned RealNVP fails to generate samples concentrated around a single class. In contrast, our neighborhood based samples capture the attributes specified by their neighborhood very well. For instance, the background, the font and the character width closely match with the neighborhood. On CIFAR-10, our model generates similar images to the specified neighbors, while the vanilla RealNVP and class conditioned RealNVP fail to generate any meaningful samples. We still observe decent coherence with neighbors using our models even though neighbors considered in this dataset are more diverse, especially in the background (due to the small dataset size of CIFAR-10). We note that the CelebA dataset comes with comprehensive attribute representations. Thus samples from attributes conditioned model tend to align with the specified attributes. However, our targeted sampling gives fine-grained control over the sample properties without the expensive (to collect) attribute labeling data. We can see that our samples capture the high-level attributes of neighbors, such as the orientation, the ornament, the hair style and the skin color.

Table 1 and Table 2 quantitatively compare the sample quality using FID and PRD respectively. We also list the FID scores of GANs for an additional comparison. We see consistent improvement over the baseline RealNVP on both recall and precision. That is, our methods utilize neighborhood as memory recall to cover diverse data space and model local manifold to better capture the data manifold. Our neighbor based targeted sampling serves to bridge the gap between non-adversarial and adversarial generative models.

Relative likelihood values can be found in Table. 3. We report bits per dimension (bpd) (Papamakarios, Pavlakou, and

Table 2: PRD scores. As in (Sajjadi et al. 2018), $(F_8, F_{\frac{1}{8}})$ pairs are reported here. GAN results are from (Sajjadi et al. 2018) (marker B, i.e., the one with best precision and recall). Best scores for likelihood methods are in **bold**.

Dataset	GAN	RNVP	CC-RNVP	NCL	NCT
MNIST	$(0.99, 0.99)$	$(0.982, 0.982)$	$(0.987, 0.973)$	$(0.981, 0.970)$	$(0.990, 0.982)$
SVHN	$(-, -)$	$(0.740, 0.731)$	$(0.785, 0.920)$	$(0.808, \mathbf{0.928})$	$(0.841, 0.841)$
CIFAR-10	$(0.77, 0.84)$	$(0.322, 0.631)$	$(0.360, 0.599)$	$(0.451, 0.647)$	$(0.568, 0.691)$
CelebA	$(0.87, 0.85)$	$(0.675, 0.766)$	$(0.835, 0.786)$	$(0.853, 0.806)$	$(0.849, 0.821)$

Table 3: Bits per dimension results for MNIST, SVHN, CIFAR-10 and CelebA. Neighborhood conditioned likelihoods $p(\cdot | \mathbf{N})$ are reported for our NCL and NCT models, while likelihoods and class conditioned likelihoods are reported for vanilla RealNVP (RNVP) and class conditioned RealNVP (CC-RNVP) respectively.

Dataset	RNVP	CC-RNVP	NCL	NCT
MNIST	0.449	0.415	0.402	0.424
SVHN	2.270	2.209	2.204	2.285
CIFAR-10	3.547	3.544	3.552	3.543
CelebA	3.018	3.150	2.931	2.934

Murray 2017) for the models considered. Note that we are not trying to compare conditional likelihoods with marginal likelihoods. We report likelihoods here only as a justification that our model is not trying to remember the training set.

Furthermore, it is interesting to note that an approximation to the marginal likelihood for our model is typically close in value to conditional neighborhood likelihoods, see appendix for details. We also explore reducing the computational cost of marginalization by using fewer neighborhoods. To do so, we cluster the training data into several larger neighborhoods. It is worth noting that our models are robust to the choice of neighborhoods, and even with far fewer neighborhoods, we observe consistent samples with the corresponding neighborhood. Please refer to appendix for more details.

3.2 Novelty Detection

Generative models have shown potential for novelty/anomaly detection, an important task in ML. Although OoD issues can potentially limit the effectiveness of likelihood based methods, we find that our neighbor based models can significantly improve detection accuracy when combined with the proposed contrastive training strategy.

Following the standard setup in (Perera, Nallapati, and Xiang 2019), we create one-class novelty detection tasks from MNIST dataset. We use the original training-testing split of MNIST dataset. Training is conducted using data of a single class from original training split. We test each model using the whole testing split. Test data from a different class the model is trained on are considered to be negative. To train our neighbor based models, we pull neighbors by Euclidean distance on PCA features, 5 neighbors are used here. Following previous works, we use Area Under the Curve (AUC) of Receiver Operating Characteristic (ROC) to measure the perfor-

mance. We threshold the conditional likelihood $p(x | \mathbf{N}(x))$ to get the ROC curve. In this experiment, we compare the vanilla RealNVP and our proposed NCL model with or without contrastive training. The margin in contrastive loss is set to 0.5 bits per dimension. We also list results of PixelCNN from (Perera, Nallapati, and Xiang 2019) for an additional comparison. Results are shown in Table. 4. We can see that the proposed contrastive training mechanism contributes a huge improvement over NCL.

4 Related Works

Density Estimation and Generative Models: Nonparametric density estimation, such as kernel density estimation, often suffers from the curse of dimensionality and does not perform well on high dimensional data like images. Recently, deep neural networks have been employed to enable flexible density estimation. (Uria, Murray, and Larochelle 2013; Uria et al. 2016; Germain et al. 2015; Gregor et al. 2013; Oord, Kalchbrenner, and Kavukcuoglu 2016; van den Oord et al. 2016) utilize neural networks to learn the conditionals factorized by the chain rule. (Dinh, Krueger, and Bengio 2014; Dinh, Sohl-Dickstein, and Bengio 2016; Kingma and Dhariwal 2018; Papamakarios, Pavlakou, and Murray 2017; Oliva et al. 2018) construct a normalizing flow based on the change of variables theorem. (Kingma and Welling 2013) proposes and optimizes a variational lower bound for the exact likelihood. GANs (Goodfellow et al. 2014) bypass explicit density estimation by adversarial training rooted in game theory. All of the aforementioned approaches try to model the whole data distribution in a single model. Our method, however, proposes to divide and conquer the density estimation using local neighborhoods. Our proposed model can potentially be integrated into all the generative models described above.

Neighbors based Generative Models: We expound on some recent approaches that make use of neighbors for generative models below. For instance, (Bansal, Sheikh, and Ramanan 2017) completes a low-resolution signal using compositions of nearest pixels in training images. (Li and Malik 2018) attempts to model a fixed code space by matching initial noisy outputs to a nearest neighbor. Note that in contrast, in addition to providing a density, our method enables targeted sampling by specifying a neighborhood.

Other Neighbor based Models: Outside of generative models, neighbor based methods have been well studied. For instance, (Weinberger and Saul 2009; Goldberger et al. 2005) apply nearest neighbors to learn a distance metric. (Boiman, Shechtman, and Irani 2008) extends classic KNN methods



Figure 3: Samples generated from MNIST, SVHN, CIFAR-10 and CelebA. Columns in order are: real images, vanilla RealNVP (RNVP) samples, class conditioned RealNVP (CC-RNVP, each column represents one class) samples, and our neighbor conditioned models’ samples. The red boxes indicate representative neighbors. The columns following each neighborhood column are samples from that neighborhood. The first column is from NCL, and the second one is from NCT.

Table 4: Novelty detection results on MNIST dataset. Performance are measured by AUC-ROC. Best scores are in **bold**.

	0	1	2	3	4	5	6	7	8	9	Mean
PixelCNN	0.531	0.995	0.476	0.517	0.739	0.542	0.592	0.789	0.340	0.662	0.618
RNVP	0.417	0.996	0.528	0.547	0.726	0.598	0.665	0.814	0.420	0.703	0.641
NCL(w/o)	0.402	0.998	0.475	0.515	0.716	0.604	0.639	0.809	0.360	0.690	0.621
NCL(w)	0.800	0.997	0.661	0.794	0.796	0.701	0.891	0.893	0.571	0.908	0.801

to a Bayes setting to perform accurate image classification. For a general discussion on nonparametric and neighborhood based methods please refer to (Wasserman 2010).

5 Conclusion

In this work, we propose multiple ways of enhancing generative models with neighborhood information. Instead of modeling the whole manifold using a single model, we divide the support into smaller neighborhoods and model the simpler local manifolds. Moreover, our approach jointly leverages the data both to learn a latent feature space and to use as neighbors to condition on a local manifold. This reduces the burden on the network capacity as the model need not memorize the manifold properties it is conditioned on. Furthermore, our approach more closely resembles human learning, which

seamlessly leverages data to learn perceptual features and to recall instances as memories.

We extend the recently proposed RealNVP and propose two neighbor conditioned RealNVP architectures to model the local distributions. The neighborhood conditioned likelihood (NCL) models the latent distribution as a Gaussian conditioned on features of the neighborhood. In contrast, neighborhood conditioned transformations (NCT) adjust the latent space based on a neighborhood.

Empirical results show that the proposed neighborhood conditioned models improve the sample quality both quantitatively and qualitatively. Our training procedure yields models with a strong coherence between samples and neighborhoods, allowing for use in targeted sampling tasks. Furthermore, these models have the potential in neighborhood interpolation

and style transfer as shown in our interpolation experiments in the appendix. Our models have the ability to generate realistic images even using far fewer neighborhoods. Combined with the contrastive training mechanism, our neighbor conditioned model significantly improves the novelty detection performance over the vanilla RealNVP.

References

- Altman, N. S. 1992. An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician* 46(3):175–185.
- Bansal, A.; Sheikh, Y.; and Ramanan, D. 2017. Pixelnn: Example-based image synthesis. *arXiv preprint arXiv:1708.05349*.
- Boiman, O.; Shechtman, E.; and Irani, M. 2008. In defense of nearest-neighbor based image classification. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, 1–8. IEEE.
- Carrier, M., and Pashler, H. 1992. The influence of retrieval on retention. *Memory & Cognition* 20(6):633–642.
- Choi, H., and Jang, E. 2018. Generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*.
- Dinh, L.; Krueger, D.; and Bengio, Y. 2014. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2016. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*.
- Germain, M.; Gregor, K.; Murray, I.; and Larochelle, H. 2015. Made: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, 881–889.
- Goldberger, J.; Hinton, G. E.; Roweis, S. T.; and Salakhutdinov, R. R. 2005. Neighbourhood components analysis. In *Advances in neural information processing systems*, 513–520.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Gregor, K.; Danihelka, I.; Mnih, A.; Blundell, C.; and Wierstra, D. 2013. Deep autoregressive networks. *arXiv preprint arXiv:1310.8499*.
- Grover, A.; Dhar, M.; and Ermon, S. 2018. Flow-gan: Combining maximum likelihood and adversarial learning in generative models. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Gutmann, M., and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 297–304.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 6626–6637.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2016. beta-vae: Learning basic visual concepts with a constrained variational framework.
- Kingma, D. P., and Dhariwal, P. 2018. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, 10236–10245.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kuhl, P. K.; Tsao, F.-M.; and Liu, H.-M. 2003. Foreign-language experience in infancy: Effects of short-term exposure and social interaction on phonetic learning. *Proceedings of the National Academy of Sciences* 100(15):9096–9101.
- Li, K., and Malik, J. 2018. Implicit maximum likelihood estimation. *arXiv preprint arXiv:1809.09087*.
- Lucic, M.; Kurach, K.; Michalski, M.; Gelly, S.; and Bousquet, O. 2018. Are gans created equal? a large-scale study. In *Advances in neural information processing systems*, 698–707.
- Nalisnick, E.; Matsukawa, A.; Teh, Y. W.; Gorur, D.; and Lakshminarayanan, B. 2018. Do deep generative models know what they don’t know? *arXiv preprint arXiv:1810.09136*.
- Oliva, J. B.; Dubey, A.; Póczos, B.; Schneider, J.; and Xing, E. P. 2018. Transformation autoregressive networks. *arXiv preprint arXiv:1801.09819*.
- Oord, A. v. d.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*.
- Papamakarios, G.; Pavlakou, T.; and Murray, I. 2017. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, 2338–2347.
- Perera, P.; Nallapati, R.; and Xiang, B. 2019. Ocgan: One-class novelty detection using gans with constrained latent representations. *arXiv preprint arXiv:1903.08550*.
- Roweis, S. T., and Saul, L. K. 2000. Nonlinear dimensionality reduction by locally linear embedding. *science* 290(5500):2323–2326.
- Sajjadi, M. S.; Bachem, O.; Lucic, M.; Bousquet, O.; and Gelly, S. 2018. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems*, 5228–5237.
- Seung, H. S., and Lee, D. D. 2000. The manifold ways of perception. *science* 290(5500):2268–2269.
- Uribe, B.; Côté, M.-A.; Gregor, K.; Murray, I.; and Larochelle, H. 2016. Neural autoregressive distribution estimation. *The Journal of Machine Learning Research* 17(1):7184–7220.
- Uribe, B.; Murray, I.; and Larochelle, H. 2013. Rnade: The real-valued neural autoregressive density-estimator. In *Advances in Neural Information Processing Systems*, 2175–2183.
- van den Oord, A.; Kalchbrenner, N.; Espeholt, L.; Vinyals, O.; Graves, A.; et al. 2016. Conditional image generation with pixelcnn decoders. In *Advances in Neural Information Processing Systems*, 4790–4798.
- Wasserman, L. 2010. *All of Nonparametric Statistics*. Springer Publishing Company, Incorporated, 1st edition.
- Weinberger, K. Q., and Saul, L. K. 2009. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research* 10(Feb):207–244.