

Normalized Diversification

Shaohui Liu^{1*} Xiao Zhang^{2*} Jianqiao Wang² Jianbo Shi²
¹Tsinghua University ²University of Pennsylvania
blueber2y@gmail.com, {zhang7, wnjq, jshi}@seas.upenn.edu

Abstract

Generating diverse yet specific data is the goal of the generative adversarial network (GAN), but it suffers from the problem of mode collapse. We introduce the concept of normalized diversity which force the model to preserve the normalized pairwise distance between the sparse samples from a latent parametric distribution and their corresponding high-dimensional outputs. The normalized diversification aims to unfold the manifold of unknown topology and non-uniform distribution, which leads to safe interpolation between valid latent variables. By alternating the maximization over the pairwise distance and updating the total distance (normalizer), we encourage the model to actively explore in the high-dimensional output space. We demonstrate that by combining the normalized diversity loss and the adversarial loss, we generate diverse data without suffering from mode collapsing. Experimental results show that our method achieves consistent improvement on unsupervised image generation, conditional image generation and hand pose estimation over strong baselines.

1. Introduction

Diversity is an important concept in many areas, e.g. portfolio analysis [33], ecological science [26] and recommendation system [45]. This concept is also crucial to generative models which have wide applications in machine learning and computer vision. Several representative examples include Variational Autoencoder (VAE) [21] and Generative Adversarial Network (GAN) [14], which are capable of modeling complicated data. One ideal principle shared by all generative models, simple or complex, is quite similar, that the generated data should be diverse. Otherwise, the model may have a so-called problem of mode collapse, where all generated outputs are highly similar. This problem is more common in GAN [14] since the objective function is mostly about the validity of generated samples but not the diversity of them.

* indicates equal contribution




Figure 1: Comparison of generative models’ capability to learn from sparse samples with unknown topology (a donut shape). Generated samples from GAN [14], BourGAN [39] and ours are illustrated. **GAN** [14] suffers from mode collapse. **BourGAN** [39] concentrates tightly around the training data points. **Ours** generates dense coverage around the training samples with limited outliers.

Our goal is to learn a generative model for high-dimensional image data which are non-uniformly distributed over a space with unknown topology, only using very sparse samples. These conditions stated above, even one of them being removed, may result in a much easier problem to solve. We formulate the solution as learning a mapping from a low-dimensional variable in latent space \mathcal{Z} to an image defined in output space \mathcal{Y} . However, we face a dilemma that, in order to fit the complicated topology of output data, we need the learned mapping to be complex and highly expressive, which actually requires dense samples from real-world. This problem might be alleviated if the topology is simpler since a parametrized function of smaller VC dimension might be able to fit. This situation however is most likely not the case if we need to deal with arbitrary images.

We start from an idea that is orthogonal to previous research: that we aim to expand as well as unfold the manifold of generated images, actively and safely. We want to model the data with complex topology in output space using a parametrized mapping from a latent space with simple topology. To learn a model that is specific to training images and also generalizable, we wish to achieve that the *interpolation* and *extrapolation* on a pair of latent variables should generate valid samples in output space. This requires the generative model, globally, to generate diverse and valid images via exploration, which we refer to as safe extrapolation, and locally, to preserve the neighbouring relation so that interpolation is also safe.

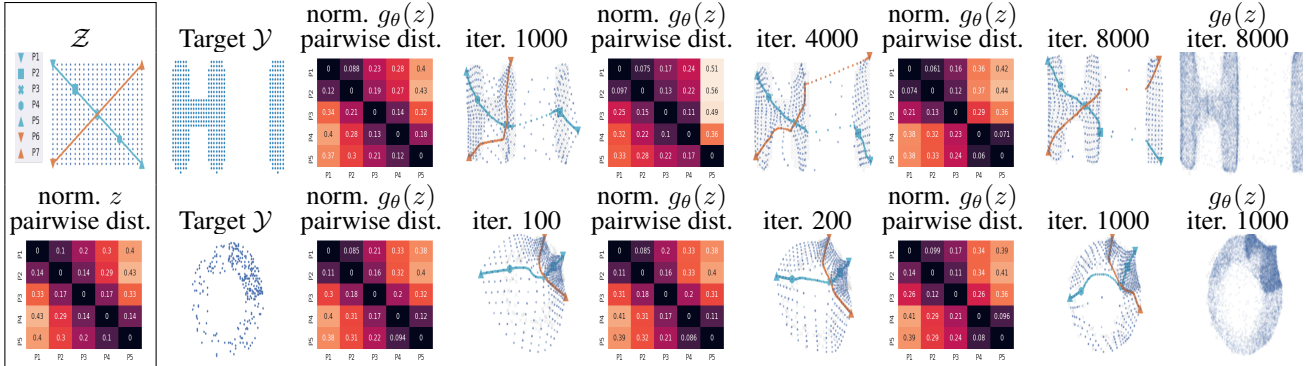


Figure 2: Illustration on the training procedure with proposed normalized diversification on highly irregular topology (top row) and non-uniform data density (bottom row). The generative model can effectively learn from sparse data by constructing a mapping from the latent space \mathcal{Z} to the target distribution by minimizing normalized diversity loss. **Top Left:** The latent variable $z \in \mathcal{Z}$ in 2D space is sampled from an uniform distribution. 5 points (P1-P5, colored blue) along the diagonal are used for illustration. **Bottom Left:** The normalized pairwise distance matrix on P1-P5. **From Left to Right:** We show qualitative results on two synthetic cases: ‘HI’ and ‘Ring’. We visualize the mapping from the latent space to the output space for several iterations together with $D^Y \in \mathbb{R}^{5 \times 5}$ and we illustrate safe interpolation on the diagonal of the latent space onto the output space. **Right Most Column:** We generate dense samples from the learned model, to illustrate the diversity and the consistency w.r.t. the ground-truth distribution shown in the second column.

In this paper, we propose *normalized diversification* following the motivation above. This is achieved by combining the adversarial loss of GAN and a novel normalized diversity loss. This diversity loss encourages the learned mapping to preserve the normalized pairwise distance between every pair of inputs $z, z' \in \mathcal{Z}$ and their corresponding outputs $g_\theta(z), g_\theta(z') \in \mathcal{Y}$, where g_θ is a parametrized function. This utilizes the same insight as manifold unfolding [37]. During training, we also fix the normalization term while maximizing the pairwise distance, which also encourages the model to visit outer modes i.e. to extrapolate.

To illustrate the concept, we sample sparse points from a donut shape as training data for GAN [14], BourGAN [39] and our method, shown in Figure 1. After the model being trained, we generate 5k new samples from the learned distribution for visualization. Our method achieves safe interpolation, fills in the gap in sparse samples and generates dense coverage while GAN [14] and BourGAN [39] fails to generalize.

Our paper is organized as follows: Section 3 presents more motivations towards a better understanding of normalized diversification. Section 4 describes how we apply normalized diversification onto multiple vision applications by employing different metric spaces on the outputs. Finally in Section 5, we show promising experimental results on multiple tasks to demonstrate the effectiveness.

2. Related Work

Early attempts for addressing the problem of mode collapse include maximum mean discrepancy [10], boundary equilibrium [4] and training multiple discriminators [9].

Later, some other works [1, 5, 27] modified the objective function. The usage of the distance matrix was implicitly introduced as a discrimination score in [32] or covariance in [20]. Recently, several novel methods via either statistical or structural information are proposed in [35, 24, 39].

Most problems in computer vision are fundamentally ill-posed that they have multiple or infinite number of solutions. To obtain a better encoder with diverse generated data, there are a variety of ideas making full use of the VAE [21], GAN [14] and its conditional version to develop better models [3, 18, 23, 31, 34, 44]. Huang et al. [17] proposed to learn a disentangled representation, Ye et al. [41] used a parametrized GMM, other works like [2, 6, 11, 12] selectively back-propagated the gradients of multiple samples. The normalized diversification, defined with pairwise terms on the mapping function itself, appears orthogonal to most of the existing methods.

3. Method

We consider a generative model g_θ that generates an image $y \in \mathcal{Y}$ based on a latent variable z . The target of training the model is to fit an unknown target distribution $p_{data}(y)$ based on limited samples. In this paper, we consider two different kinds of implicit generative models that solve different problems but intrinsically share the same spirit of diversification.

- *Unsupervised generative model.* This model is used for tasks that do not depend on auxiliary information, e.g. image synthesis without supervision. Existing methods such as GAN [14] and VAE [21] use a latent

variable $z \in \mathcal{Z}$ that follows a parametric distribution $p(z)$ in latent space and learn the mapping $g_\theta : \mathcal{Z} \rightarrow \mathcal{Y}$ to fit the target $y \in \mathcal{Y}$.

- *Conditional generative model.* This model utilizes additional information to generate more specific outputs, e.g. text-to-image (image-to-image) translation, pose estimation, future prediction. Related works include a variety of conditional generative models [3, 18, 23, 31, 44]. They also use a predefined latent space \mathcal{Z} and aim to fit the joint distribution $(\mathcal{X}, \mathcal{Y})$ for the input domain \mathcal{X} and output domain \mathcal{Y} (for convenience we also use \mathcal{Y} here). Specifically, an encoder $E : \mathcal{X} \rightarrow \mathcal{C}$ is used to get the latent code $c \in \mathcal{C}$. Then, either addition (VAE) or concatenation (CGAN) is employed to combine c and z for training generator $g_\theta : \mathcal{C} \times \mathcal{Z} \rightarrow \mathcal{Y}$.

The problem of mode collapse is frequently encountered especially in generative models like GAN [14], that the model generates highly similar data but satisfies the criterion for training. Our motivation for normalized diversification is to encourage g_θ to generate data with enough diversity, so the model visits most of the important modes. In the meantime, we wish g_θ to be well-conditioned around the visited modes, so we could infer latent variables from some valid image samples, and ensure safe interpolation and extrapolation between these latent variables to generate meaningful images. We could address the problem of mode collapse by diversifying the outputs, which is similar to enlarging the variance parameter of a Gaussian distribution, but a hard problem is how to measure the diversity of real-world images analytically and properly, as the variance of a distribution might be too universal for general tasks to be specific to our problem.

We start with an intuition to prevent the pairwise distance $d(\cdot, \cdot)$ between two generated points A and B from being too close. Note that if the outputs are linearly scaled, so will the pairwise distance $d(A, B)$: the samples seem to diversify from each other, without actually solving the intrinsic problem of mode collapse. Thus, we measure whether the mapping preserves the normalized pairwise distance between inputs $\|z - z'\|$ and outputs $\|g_\theta(z) - g_\theta(z')\|$.

Since we do not have access to infinite amount of data, we sample a limited amount of data from a well-defined parametric distribution $p(z)$, and try to visit more mode via *diversification*. We measure the diversity of generated samples by the pairwise distance in \mathcal{Y} through the parametrized mapping g_θ . The objective function is simplified as a finite-sum form on N samples $\{z_i\}_{i=1}^N$, along with corresponding images $\{y_i | y_i = g_\theta(z_i)\}_{i=1}^N$. We denote two metric spaces $M_Z = (\mathcal{Z}, d_Z)$ and $M_Y = (\mathcal{Y}, d_Y)$. We use two additional functions h_Y and h_Z for some task-specific usage, and then define the metric as composite functions upon Eu-




Figure 3: Quantitative comparison between normalized diversity loss and BourGAN loss [39] on the learned distribution in Fig. 1. We discretized the donut region into uniform mesh grids and measured “cover rate” (the percentage of grids which have generated samples in them). We also measured “outlier rate”: the ratio of samples outside the donut. “Data sparsity” measures the cover rate of the training samples over the donut. Our method improve the cover rate over GAN [14] and BourGAN [39] while maintaining comparable outlier rate.

clidean distance, as follows,

$$d_Z(z_i, z_j) = \|h_Z(z_i) - h_Z(z_j)\|_2 \quad (1)$$

$$d_Y(y_i, y_j) = \|h_Y(y_i) - h_Y(y_j)\|_2 \quad (2)$$

To explain with a concrete application in diverse video synthesis that we try to generate a realistic video from a moving vehicle based on a sequence of segmentation mask, here d_Y may be used for extracting deep features using an off-the-shelf network like the perceptual loss [19]. We also define $D^Z, D^Y \in \mathbb{R}^{N \times N}$ to be the normalized pairwise distance matrix of $\{z_i\}_{i=1}^N$ and $\{y_i\}_{i=1}^N$ respectively in Eq.(3). Each element in the matrix is defined as

$$D_{ij}^Z = \frac{d_Z(z_i, z_j)}{\sum_j d_Z(z_i, z_j)}, \quad D_{ij}^Y = \frac{d_Y(y_i, y_j)}{\sum_j d_Y(y_i, y_j)}, \quad (3)$$

for $\forall i, j \in [N]$. The normalized diversity loss can be formulated as in Eq.(4), where $\alpha \approx 1$ is a slack factor of this diversity loss.

$$\mathcal{L}_{div}(\theta, p) = \frac{1}{N^2 - N} \sum_{i=1}^N \sum_{j \neq i}^N \max(\alpha D_{ij}^Z - D_{ij}^Y, 0). \quad (4)$$

To enforce the extrapolation, we treat the normalizer in Eq.(3) as a constant when back-propagating the gradient to generator. As a result minimizing Eq.(4) would also force the model to actively expand in the high-dimensional output space, we refer to Algorithm 1 and Section 3.1 for more details.

For the purpose of unfolding the manifold, we could focus on the expansion of densely connected pairs of $D_{ij}^Z > D_{ij}^Y$ or on the contraction of loosely connected pairs of $D_{ij}^Z < D_{ij}^Y$. Only minimizing $\max(\alpha D_{ij}^Z - D_{ij}^Y, 0)$ would encourage active extrapolation when we hold the normalizer constant.

Combining the diversity loss, the objective function over θ can be written in a compact form as

$$\min_{\theta} \mathcal{L}(\theta, p) = \mathcal{L}_{gen}(\theta, p) + \mathcal{L}_{div}(\theta, p) \quad (5)$$

where $\mathcal{L}_{gen}(\theta, p)$ is the original objective function for learning the generator g_{θ} . In the conditional model, the objective function also depends on x which we omitted here.

3.1. Interpolation and Extrapolation

A fundamental motivation behind the normalized diversification is to achieve that *interpolation* and *extrapolation* on a pair of latent variables should generate valid samples in output space. This pursues a trade-off between pairwise diversity and validity. This motivation aligns with the concept of *local isometry* [38], which has a nice property that the near neighbors in \mathcal{Z} , are encoded to \mathcal{Y} using a rotation plus a translation. Local isometry requires the output manifold to be smooth and without invalid ‘holes’ within the neighbor region of a valid point, and the interpolation in \mathcal{Z} generates valid points in \mathcal{Y} through a parametrized mapping g_{θ} where θ is the learned parameter. Differently, Variational Autoencoder (VAE) [21] mostly cares about local perturbation, or pointwise diversity.

Normalized diversification is orthogonal to previous research as it aims to unfold the manifold of generated images [37]. **Interpolation.** By pushing apart the pairwise distance of the sample points, we prevent the ‘short-cuts’ that links samples through the exterior space. As shown in Figure 1, given a set of points, our goal is to discover the underlying parametrization so we can densely generate new valid samples in the interior (on the donut), without crossing over to the exterior (the donut hole). This leads to safe interpolation. **Extrapolation.** For active exploration of output space \mathcal{Y} with the current model, in each iteration, we first calculate the normalizer of the pairwise distance matrix, then use the gradient back-propagated from the $d_{ij}^{\mathcal{Y}}$ to force expansion, after which we update the normalizer. With these alternating steps, we ensure the stability of exploration. We illustrate the evolution of the training procedure on two synthetic 2D distributions in Figure 2

3.2. Understanding from Geometric Perspective

Simply trying to enlarge the pairwise distance between samples in \mathcal{Y} can explore the unobserved space, but a crucial problem is how to make sure the interpolated points are still valid. From a geometric perspective, imagine that A, B, C are on a curvy 1D line in 2D space, the transitive distance between them, $d(A, B) + d(B, C)$ on the 1D curve is much longer than 2D distance $d(A, C)$, which violates the triangle inequality. If we make an interpolation point D in the inner part of the line (A, C) , although the direct 2D distance between $d(A, C)$ could be very small, the line

between them might lie in the part of unreasonable space out of the manifold. However, by pushing (A, C) as far as possible away, we ‘discover’ the true 1D distance.

This insight is different from existing approaches, e.g. BourGAN [39], which aims at matching the pairwise distance between output space and latent space which can preserve the data modality but hinder its generalization ability. Our method, besides the mode preserving ability, can also actively expand and unfold the manifold which enables safe interpolation on the latent space where the generated samples will neither lie outside the valid region nor overfit the existing data (See Figure 1 and 3).

3.3. Understanding via Simplified Models

To understand the normalized diversification, we start from another perspective by using simple functions to illustrate the functionality of the regularization term. We assume the generator to be a simple linear model as $g_{\theta}(z) = \theta^T z$, where θ is the matrix that characterizes the linear transformation. A simple calculation induces that the diversity regularization thus encourages $\|\theta\|_{*}$ (or $\|\theta\|_2$) being sufficiently large. However, this alone does not prevent some degenerated cases. Suppose that the generator is constructed in the following way

$$g_{\theta}(z) = \theta^T z, \quad \theta = U^T \text{diag}[\beta, 0, 0, \dots, 0] V \in \mathbb{R}^{K \times D},$$

where K and D are the dimension of the latent variable z and generated data respectively, and U, V are the matrices consisting of all singular-vectors. When β is sufficiently large, this model seems to be strongly diversified, however, is actually not a reasonable model as it measures the difference of two vectors along only one direction, i.e. the singular-vector corresponding to the singular value β . Normalization over the diversity helps to prevent these kinds of degenerated cases, as the normalized distance does not scale with β . The normalization also helps to adapt to a conditional setting that some related works e.g. BourGAN [39] might fail to adapt, based on an input variable x and a latent variable z , especially with large variations w.r.t. x . Imagine a simple example where $y = g_{\theta}(x, z) = (x + z)^3$, and $x \in [1, 10]$. The upper bound of Lipschitz constant could be larger than $3\max(x)^2 = 300$ while the lower bound should be lower than $3\min(x)^2 = 3$. However, this value is meaningless for the most part of the whole domain $[1, 10]$. Normalizing the distance could fix it to a reasonable range.

Another side problem is on how to keep the semi-definiteness of the distance metric, to concord with metric learning research [8, 37, 38, 40]. Under this property, a pair of samples (z, z') deviating from each other in whichever direction, should increase the pairwise distance in the output space. This property should also be enforced by diversification, although may not explicitly. This formulation of taking nonnegative part using $\max(\cdot, 0)$ encourages the




Figure 4: Qualitative results on ‘2D Gaussian ring’ and ‘2D Gaussian grid’. For baseline methods, We directly used the visualization results in [39].

learned metric in \mathcal{Y} to be positive semidefinite (PSD). To understand with a counter-example, we consider a metric defined as

$$d_Y(y, y') = (y - y')^T \Phi (y - y'), \quad \text{where } \Phi \in \mathbb{R}^{D \times D}.$$

We assume the eigen-decomposition of Φ to be

$$\Phi = U^T \text{diag}[-\beta, 100\beta, 0, 0, \dots, 0]U \quad (6)$$

where $\beta > 0$ (eigen-decomposition), and Φ is not PSD. Suppose there is a pair of latent variables (y, y') that $y - y' = \sum_d U_d / D$, where summation is taken over all eigenvectors. This pair of data incurs a variation of 99β of \mathcal{L}_{ndiv} in Eq.(4) if the nonnegative operator $\max(\cdot, 0)$ is not used, which leads to the objective function being minimized if β positively scales, but clearly, the metric is not reasonable as it is not PSD. However, if the operator is applied, the same case incurs a variation of β of \mathcal{L}_{ndiv} from a deviation in the subspace corresponding to the negative eigenvalue $-\beta$, as the $\max(\cdot, 0)$ operator strongly contrast those directions like U_1 that wrongly contribute to \mathcal{L}_{ndiv} , and ignore other subspace even if they correspond to strongly positive eigenvalues.

4. Real-World Applicability

In this section, we present the real-world applicability of normalized diversification for different vision applications. For different tasks, we need a validity checking function F of the generated samples to validate that the data satisfies some domain-specific constraints. For example, for the unsupervised setting, F can be interpreted as the trained discriminator in GAN [14].

Our training pipeline is summarized in Algorithm 1. For a given latent distribution $p(z)$ and the target distribution p_{data} , we sample a finite set of latent samples $\{z_i\}_{i=1}^n$ from $p(z)$ and $\{y_i\}_{i=1}^n$ from p_{data} . We then compute normalized pairwise distance matrix over $\{z_i\}_{i=1}^n$ and the generated samples $\{g_\theta(z_i)\}_{i=1}^n$ with the functions d_Z and d_Y . We update the normalizer and block its gradients as previously

Algorithm 1 Training generative model with normalized diversity.

```

1: Given: Latent distribution  $p(z)$ , Target distribution  $p_{data}$ 
2: Given: Current generator  $g_\theta$ , Distance function  $d_Z, d_Y$ 
3: function NORMDIST( $\{q_j\}_{j=1}^N, d$ ):
4:   for  $i \leftarrow 1$  to  $N$  do
5:      $S_i = \sum_{j=1}^N d(q_i, q_j)$ 
6:      $S_i \leftarrow S_i.value$  (treated as constant value with
       back-propagated gradients blocked)
7:      $D_{ij} = d(q_i, q_j) / S_i, j \in [N]$ 
8:   end for
9:   return  $D$ 
10: end function
11: while not converged do
12:    $\{z_i\}_{i=1}^N \sim p(z)$ 
13:    $\{y_i\}_{i=1}^N \sim p_{data}$ 
14:    $D^Z = \text{NORMDIST}(\{z_i\}_{i=1}^N, d_Z)$ 
15:    $D^Y = \text{NORMDIST}(\{g_\theta(z_i)\}_{i=1}^N, d_Y)$ 
16:   (Optional for GAN) Update the discriminator  $F$ .
17:   Compute total loss  $\mathcal{L}$  by Eq. 5
18:   Update model parameter  $\theta = \theta - \eta \partial \mathcal{L} / \partial \theta$ 
19: end while

```

discussed. Finally, we compute the total loss \mathcal{L} in Eq.(5) and update the model with the back-propagated gradients for each step.

4.1. Applications

The proposed normalized diversity term is general and can be applied to many vision applications including image generation, text-to-image (image-to-image) translation and hand pose estimation, etc. We specify the conditional input domain \mathcal{X} , the output domain \mathcal{Y} and the function h_Y for each application to compute normalized diversity loss in Eq.(4). We use a predefined \mathcal{Z} with uniform distribution and $h_Z(z) = z$ for all tasks.

- **Image generation.** There is no conditional input for this unsupervised generative model setting. \mathcal{Y} is the output image domain. The normalized pairwise distance on \mathcal{Y} can be computed either with ℓ_1 or ℓ_2 distance or employing deep metrics such as using the GAN [14] discriminator as the function h_Y .
- **Conditional image generation.** The conditional input domain \mathcal{X} can be the input text and input image for text-to-image and image-to-image translation respectively. \mathcal{Y} is the output image domain. The computation of normalized pairwise distance on \mathcal{Y} is similar to that in image generation.
- **Hand pose estimation.** The conditional input domain \mathcal{X} is the input RGB hand image. $\mathcal{Y} \in \mathbb{R}^{3K}$ is the

Table 1: Results on the synthetic dataset. We followed the experimental setting in [39].

Method	2D Ring		2D Grid	
	#modes	fail (%)	#modes	fail (%)
GAN [14]	1.0	0.1	17.7	17.7
Unrolled [27]	7.6	12.0	14.9	95.1
Ours	8.0	10.3	25.0	11.1
VEEGAN [35]	8.0	13.2	24.4	22.8
PacGAN [24]	7.8	1.8	24.3	20.5
BourGAN [39]	8.0	0.1	25.0	4.1

output 3D pose, where K denotes the number of the joints. For the function h_Y , we employ the visibility mask $V \in \{0, 1\}^{3K}$ to gate each joint of the output 3D pose inversely, encouraging diversity on the occluded joints, i.e. $h_Y(y) = (1 - V) \circ y$.

Note that [39] cannot fit in the conditional setting because the number of possible conditional inputs, e.g. input text, RGB hand image are infinite, making the pairwise information for each conditional input impossible to be pre-computed. Specifically for hand pose estimation, we develop a pseudo-renderer by using morphological clues to get a visibility mask.

The application of normalized diversity can fit in arbitrary checking function F for different tasks. When compared to baseline adversarial methods, we use adversarial loss for image generation either unsupervised or conditional. For hand pose estimation, we choose to combine the l_2 distance on the visible joints and the joint adversarial discriminator of (x_i, \hat{y}_i) (image-pose GAN), where \hat{y}_i denotes the 2D projection of the output 3D pose y_i .

$$F(x, y) = F_{vis}(x, y) + F_{gan}(x, y) \quad (7)$$

The two checking functions are formulated as follows

$$F_{vis}(x, y) = \|V \circ (y - y_r)\|_2^2 \quad (8)$$

$$F_{gan}(x, y) = \mathbb{E}_{x, \hat{y}_r \sim p_{data}} [\log D(x, \hat{y}_r)] \quad (9)$$

$$+ \mathbb{E}_{x \sim p_{data}, z \sim p(z)} [\log(1 - D(x, \hat{y}))]$$

With the image-pose GAN, the system has the capability to learn much subtle relationship between x and y . By constraining only on the visible joints, our system gets less noisy gradients and can learn better pose estimation along with better image features.

5. Experiments

We conducted experiments on multiple vision tasks under both unsupervised setting (Section 5.1 and 5.2) and conditional setting (Section 5.3 and 5.4) to demonstrate the effectiveness of the proposed idea.

Table 2: FID Results [16] on Image Synthesis. ‘‘SN’’ and ‘‘GP’’ denote spectral normalization [28] and gradient penalty [15] respectively. Our method achieved consistent improvement over various GANs.

	CIFAR-10		CelebA	
	w/o ndiv	w ndiv	w/o ndiv	w ndiv
GAN+SN	23.7	22.9	10.5	10.2
GAN+SN+GP	22.9	22.0	9.4	9.1
WGAN+GP	25.1	23.9	9.9	9.5
WGAN+SN+GP	23.7	23.3	9.2	9.0

Compared to the baseline, we added the normalized diversification loss with all other settings unchanged. We simply used $\alpha = 0.5$ for Eq.(4) in most experiments except conditional image generation, where we used $\alpha = 0.8$.

5.1. Synthetic Datasets

We tested our method on the widely used 2D Gaussian ring and 2D Gaussian grid to study the behavior of mode collapse. For each iteration in our method, we performed line 16 and 18 in Algorithm 1 all repeatedly for 3 times for faster convergence. Results are shown in Table 1. We compared several baselines [14, 27, 35, 24, 39]. Of those methods VEEGAN [35], PacGAN [24] and BourGAN [39] require pairwise information of the real image sets, which is an expensive pre-requisite and cannot fit in the conditional setting. Moreover, BourGAN [39] suffered from severely overfitting the training samples (See Fig. 1). Our method, with the normalized diversification, achieved promising results in the long-term training (50k iterations) with only on-demand pairwise information. In Fig. 4, we qualitatively demonstrate the results on synthetic point sets where our method achieve comparable or even better results over VEEGAN [35] and PacGAN [24] without using two or more real samples at a time.

5.2. Image Generation

We further tested our method on the widely accepted task of image generation. We conducted experiments on CIFAR-10 [22] and CelebA [25] with state-of-the-art methods via the off-the-shelf library¹. Because there exist relatively dense samples in this setting, conventional methods fit the problem relatively well. Nevertheless, as shown in Table 2, with the normalized diversification loss, our method achieves consistent improvement on CIFAR-10 [22] and CelebA [25] over strong baselines with spectral normalization [28] and gradient penalty [15]. By regularizing the model behavior around the visited modes through normalized diversification, the generator maintains a safe extrapolation depending on the current outputs.

¹https://github.com/google/compare_gan

Table 3: Results of text-to-image translation on [30]. 15 volunteers voted for the better of the randomly sampled pairs for a random text on the webpage. Diversity was computed via the perceptual metric [43] using Inception model [36].

	votes	diversity
GAN-CLS [31]	287 (31.3%)	43.7±16.8
Ours	291 (31.7%)	58.1±18.5
Neutral	340 (37.0%)	-

Table 4: Quantitative comparison with BicycleGAN [44] on conditional facade image generation [7]. We computed FID [16] with the real data to measure the image quality (lower better) and used LPIPS [43] to measure the output diversity (higher better).

	FID ↓	LPIPS ↑
Real data	0.0	0.291
BicycleGAN [44]	83.0	0.136 ± 0.049
Ours	76.3	0.188 ± 0.048

5.3. Conditional Image Generation

We tested the performance of conditional image generation described in Section 4.1. For text-to-image translation, we directly followed the architecture used in [31] for fair comparison and added the diversity loss term to the overall objective function. We used a discriminator as h_Y to extract features from synthesized images and measured the pairwise image distance in feature space. Experiments were conducted on Oxford-102 Flower dataset [30] with 5 human-generated captions per image. We generated 5 samples for each text at training stage. Both the baseline and our method were trained for 100 epochs on the training set. We employed the zero-shot setting where we used the 3979 test sets to evaluate both methods. We evaluated the quality and the diversity of the generated images. For the diversity measurement, we first generated 10 images for each input text, then used the perceptual metric [43] via Inception model [36] to compute the variance of the 10 generated samples and took average over test sets. Table 3 shows the results on [30]. With comparable image quality, we significantly improved upon the baseline in terms of the diversity.

We compare with a strong baseline, BicycleGAN [44], on image-to-image translation to further demonstrate our ability to model multimodal distribution. We removed the conditional latent regression branch of BicycleGAN [44] and added normalized diversification. We used ℓ_1 as the pairwise image distance. Table 4 and Fig. 5 show the results on [7]. Even without the conditional latent regression branch [44], our method outperforms BicycleGAN [44] on both image quality and diversity.




Figure 5: Qualitative results of conditional image generation on facade dataset [7]. Our method improves both image quality and diversity over BicycleGAN [44]. **Top left:** input image. **Bottom Left:** corresponding groundtruth image. **Top Right:** five generated images from BicycleGAN [44]. **Bottom Right:** five generated images from Ours.

5.4. Hand Pose Estimation

We conducted experiments on three RGB hand datasets including GANerated hands [29], Stereo [42], and FPHAB [13]. For all datasets, we manually cropped the padded hand bounding box and resized the input to 128x128. Following [29], we directly used the net architecture of their released model. The weights were initialized from ImageNet pretrained model. For each image, our multimodal system generated 20 samples for training and 100 samples for testing. We used $z_{dim} = 10$ for all the conditional model. Our method used channel-wise concatenation to aggregate the encoded features c_i and z_i . Each dimension of z_i was sampled from a uniform distribution $U(0, 1)$.

Evaluating the multimodal predictions is a non-trivial task. Conventional methods put a max operation on top of the multiple predictions and use the samples nearest the groundtruth for evaluation. Some works even use this max operation over each joint. This results in relatively unfair comparison because simply drawing a sample uniformly distributed in $U(0, 1)$ will result in near zero error. Thus, we introduce a better evaluation protocol for the multimodal hand pose estimation: 1) Visible joint accuracy. For the visible joints, we use the Percentage of Correct Keypoints (PCK) following conventional methods. 2) Standard deviation. We compute standard deviation of the outputs for each

Table 5: Visible joint accuracy on GANerated hands [29]. Using a conditional model along with the proposed normalized diversification helps on the visible joints predictions. Our method achieves 1.6 (91.6→93.2) accuracy gain over deterministic regression.

Regression	91.6
Ours w/o diversity loss	91.7
Ours	92.8
Ours+	93.2

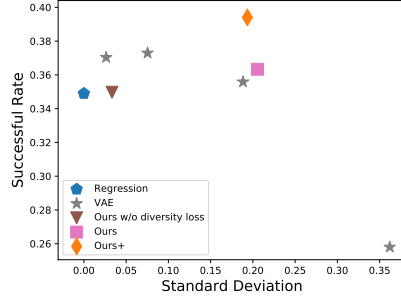


Figure 6: Results on successful rate and standard deviation. We get significant improvements on both quality and diversity.

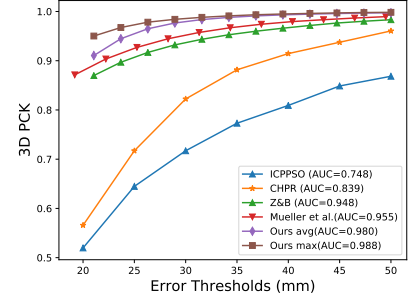


Figure 7: Results on Stereo dataset [42]. Our method outperforms existing state-of-the-art alternatives by a large margin.

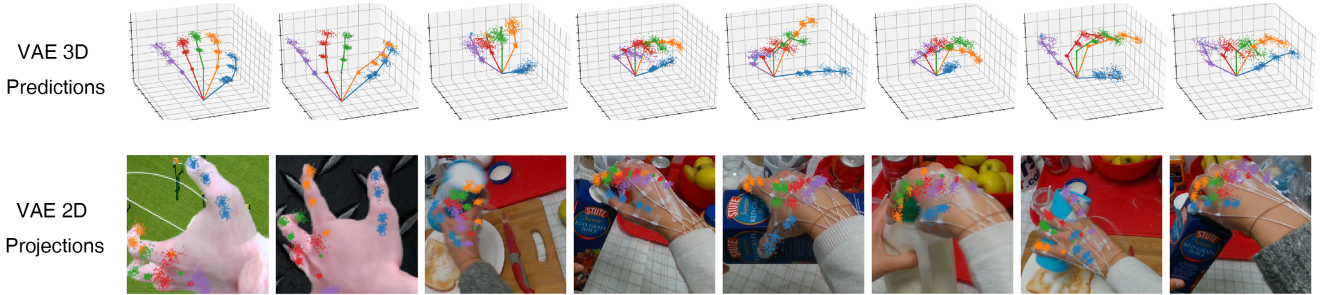


Figure 8: Qualitative results of the multiple pose predictions on GANerated hands [29] and FPHAB [13]. We show 3D hand predictions and its projections on 2D image (better viewed when zoomed in). With comparable variance, while VAE [21] failed to generate high quality samples, our model generates multiple valid 3D poses with subtle image-pose correlations.

image and take average. 3) Successful rate. For each image, we use the pre-computed hand mask and object mask to check validity of the samples. The sample is considered valid if the whole hand configuration including the occluded joints lie inside the foreground mask.

We first compare our method with the deterministic regression method on the GANerated hand dataset [29]. We concatenate the encoded feature $c_i \in \mathbb{R}^{100}$ with z_i at bottleneck which is contrast to the architecture of VAE [21] where they restricted the latent space z_i to be in the form of Gaussian distribution. ‘Ours+’ denotes our model adapting a higher dimensional latent space $c_i \in \mathbb{R}^{16384}$. We use 100 dimensional latent space z_i in all experiments. Combining Table 5 and Figure 6, it is clear that using a one-to-many conditional model benefits much on the visible joints as well as the successful rate. Our method achieves significant improvements on all three metrics, while VAE [21] struggles more on its quality-diversity trade-off.

Moreover, we tested our method on two real-world datasets including the Stereo Tracking Benchmark [42] and FPHAB [13]. Our method outperforms state-of-the-art methods both quantitatively and qualitatively. As is shown in Figure 8, our method captures much subtle ambiguity and could generate accurate predictions for visible joints on

each sample. With normalized diversification, the model maintains a centralized structure which has good property for pairwise interpolation. This promotes “safer” extrapolation for robust occluded joints detection which leads to valid yet diversified outputs.

6. Conclusion

In this paper, we proposed *normalized diversification*, a generalized loss on the mapping function measuring whether the mapping preserves the relative pairwise distance to address the problem of mode collapse. We aim to diversify the outputs with normalized pairwise distance, encouraging safe interpolation in the latent space and active extrapolation towards outer important states simultaneously. Results show that by employing different metric spaces, normalized diversification can be applied to multiple vision applications and achieves consistent improvements on both encoding quality and output diversity.

Acknowledgements

We gratefully appreciate the support from Honda Research Institute Curious Minded Machine Program. We sincerely thank Dr. Gedas Bertasius and Xin Yuan for valuable discussions.

References

- [1] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, pages 214–223, 2017. 2
- [2] A. Bansal, Y. Sheikh, and D. Ramanan. Pixelnn: Example-based image synthesis. 2018. 2
- [3] J. Bao, D. Chen, F. Wen, H. Li, and G. Hua. Cvae-gan: Fine-grained image generation through asymmetric training. In *ICCV*, pages 2745–2754, 2017. 2, 3
- [4] D. Berthelot, T. Schumm, and L. Metz. Began: boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717*, 2017. 2
- [5] T. Che, Y. Li, A. P. Jacob, Y. Bengio, and W. Li. Mode regularized generative adversarial networks. In *ICLR*, 2017. 2
- [6] Q. Chen and V. Koltun. Photographic image synthesis with cascaded refinement networks. In *ICCV*, volume 1, page 3, 2017. 2
- [7] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, pages 3213–3223, 2016. 7
- [8] J. V. Davis, B. Kulis, P. Jain, S. Sra, and I. S. Dhillon. Information-theoretic metric learning. In *ICML*, pages 209–216. ACM, 2007. 4
- [9] I. Durugkar, I. Gemp, and S. Mahadevan. Generative multi-adversarial networks. In *ICLR*, 2017. 2
- [10] G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. In *UAI*, 2015. 2
- [11] H. Fan, H. Su, and L. J. Guibas. A point set generation network for 3d object reconstruction from a single image. In *CVPR*, volume 2, page 6, 2017. 2
- [12] M. Firman, N. D. Campbell, L. Agapito, and G. J. Brostow. Diversenet: When one right answer is not enough. In *CVPR*, pages 5598–5607, 2018. 2
- [13] G. Garcia-Hernando, S. Yuan, S. Baek, and T.-K. Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *CVPR*, 2018. 7, 8
- [14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014. 1, 2, 3, 5, 6
- [15] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. In *NIPS*, pages 5767–5777, 2017. 6
- [16] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, pages 6629–6640. 2017. 6, 7
- [17] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018. 2
- [18] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 2, 3
- [19] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 3
- [20] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018. 2
- [21] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1, 2, 4, 8
- [22] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009. 6
- [23] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther. Autoencoding beyond pixels using a learned similarity metric. In *ICML*, pages 1558–1566, 2016. 2, 3
- [24] Z. Lin, A. Khetan, G. Fanti, and S. Oh. Pacgan: The power of two samples in generative adversarial networks. In *NeurIPS*, 2018. 2, 6
- [25] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738, 2015. 6
- [26] A. E. Magurran. Why diversity? In *Ecological diversity and its measurement*, pages 1–5. Springer, 1988. 1
- [27] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein. Unrolled generative adversarial networks. In *ICLR*, 2017. 2, 6
- [28] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018. 6
- [29] F. Mueller, F. Bernard, O. Sotnychenko, D. Mehta, S. Sridhar, D. Casas, and C. Theobalt. Generated hands for real-time 3d hand tracking from monocular rgb. In *CVPR*, pages 49–59, 2018. 7, 8
- [30] M.-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Computer Vision, Graphics & Image Processing, 2008. ICVGIP'08. Sixth Indian Conference on*, pages 722–729. IEEE, 2008. 7
- [31] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. In *ICML*, pages 1060–1069, 2016. 2, 3, 7
- [32] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen. Improved techniques for training gans. In *NIPS*, pages 2234–2242. 2016. 2
- [33] B. H. Solnik. Why not diversify internationally rather than domestically? *Financial analysts journal*, pages 48–54, 1974. 1
- [34] A. Spurr, J. Song, S. Park, and O. Hilliges. Cross-modal deep variational hand pose estimation. In *CVPR*, pages 89–98, 2018. 2
- [35] A. Srivastava, L. Valkoz, C. Russell, M. U. Gutmann, and C. Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. In *NIPS*, pages 3308–3318, 2017. 2, 6
- [36] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 7
- [37] K. Q. Weinberger, J. Blitzer, and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. In *NIPS*, pages 1473–1480, 2006. 2, 4

- [38] K. Q. Weinberger and L. K. Saul. Unsupervised learning of image manifolds by semidefinite programming. *IJCV*, 70(1):77–90, 2006. [4](#)
- [39] C. Xiao, P. Zhong, and C. Zheng. Bourgan: Generative networks with metric embeddings. In *NeurIPS*, 2018. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [40] E. P. Xing, M. I. Jordan, S. J. Russell, and A. Y. Ng. Distance metric learning with application to clustering with side-information. In *NIPS*, pages 521–528, 2003. [4](#)
- [41] Q. Ye and T.-K. Kim. Occlusion-aware hand pose estimation using hierarchical mixture density network. In *ECCV*, 2018. [2](#)
- [42] J. Zhang, J. Jiao, M. Chen, L. Qu, X. Xu, and Q. Yang. 3d hand pose tracking and estimation using stereo matching. *arXiv preprint arXiv:1610.07214*, 2016. [7](#), [8](#)
- [43] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [7](#)
- [44] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman. Toward multimodal image-to-image translation. In *NIPS*, pages 465–476, 2017. [2](#), [3](#), [7](#)
- [45] C.-N. Ziegler, S. M. McNee, J. A. Konstan, and G. Lausen. Improving recommendation lists through topic diversification. In *WWW*, pages 22–32. ACM, 2005. [1](#)