# A GEOMETRIC UNIFICATION OF GENERATIVE AI WITH MANIFOLD-PROBABILISTIC PROJECTION MODELS

## Anonymous authors

000

001

002

004 005 006

008 009 010

011

013

014

015

016

017

018

019

021

022

025

026

027

028

029

031

032 033 034

037

039

042 043 044

045

046

048 049

051 052 Paper under double-blind review

## Abstract

The foundational premise of generative AI for images is the assumption that images are inherently low-dimensional objects embedded within a highdimensional space. Additionally, it is often implicitly assumed that thematic image datasets form smooth or piecewise smooth manifolds. Common approaches overlook the geometric structure and focus solely on probabilistic methods, approximating the probability distribution through universal approximation techniques such as the kernel method. In some generative models the low dimensional nature of the data manifest itself by the introduction of a lower dimensional latent space. Yet, the probability distribution in the latent or the manifold's coordinate space is considered uninteresting and is predefined or considered uniform. This study unifies the geometric and probabilistic perspectives by providing a geometric framework and a kernel-based probabilistic method simultaneously. The resulting framework demystifies diffusion models by interpreting them as a projection mechanism onto the manifold of "good images". This interpretation leads to the construction of a new deterministic model, the Manifold-Probabilistic Projection Model (MPPM), which operates in both the representation (pixel) space and the latent space. We demonstrate that the Latent MPPM (LMPPM) outperforms the Latent Diffusion Model (LDM) across various datasets, achieving superior results in terms of image restoration and generation.

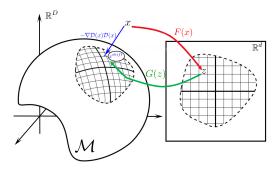


Figure 1: Illustration of our manifold-aware restoration approach. The blue path shows direct projection onto manifold  $\mathcal{M}$  using distance function  $\mathcal{D}_{\mathcal{M}}(x)$ , while the red-green path represents encoding-decoding through latent space  $\mathbb{R}^d$  via functions F and G. Ideally, both paths converge to the same manifold point, ensuring geometrically consistent restoration.

#### 1 Introduction

Restoration of images refers to the inverse process of generating a clean, meaningful, and non-corrupted image from a noisy, blurred, or otherwise degraded input. A critical aspect of

this process involves the use of prior knowledge or a well-approximated distribution function over the set of clean images within a specific class. In this work, we propose the manifold assumption, which asserts that the set of desired images resides on a low-dimensional smooth manifold. We integrate this assumption with a probabilistic perspective. Specifically, we extend the conventional Monge patch description of the data manifold, typically provided by generative models such as autoencoders (AE) (Rumelhart & McClelland, 1987), variational autoencoders (VAE) (Kingma & Welling, 2013), and generative adversarial networks (GAN) (Goodfellow et al., 2014). Our approach augments this description by introducing a distance function that assigns, for each point in the pixel (ambient/representation) space, the distance to the closest point on the manifold. We treat here images as primary examples, but evidently it can be applied to any dataset that has this manifold structure. Next, we establish a connection between the geometric framework and the probabilistic perspective by introducing a geometric-based probability function and its kernel-based approximation. We further relate these approaches to diffusion-like methods, utilizing the score function to generate, in the ambient space, a vector field that directs each noisy or corrupted image towards the closest point on the manifold of clean images. By iteratively following this vector field, a diffusion-like flow is generated, guiding the corrupted image progressively towards a clean image residing on the manifold.

To accommodate the possibility of a nonuniform probability distribution on the manifold, we employ a kernel method that adjusts the diffusion-like flow to balance the trade-off between proximity to the manifold and the probability of a point on the manifold representing a clean and meaningful image. This integration of geometric principles with the kernel method constitutes the primary novelty of our approach. Furthermore, we extend these general concepts, the distance function, score, and diffusion-like flow, to operate within the latent space, thereby reducing computational complexity and enhancing the accuracy of the distance function. We evaluated our proposed method on the MNIST and SCUT-FBP5500 datasets, demonstrating superior performance compared to a leading method such as the Latent Diffusion Model (LDM) (Rombach et al., 2022).

#### 1.1 Related Work

054

055

057

058

059

060

061

062

063

064

065

066

067

068

069

070 071

072

073

074

075

076

077

078

079

080 081 082

083

085

086

087

880

090

091

092

096

098

100

101

102

103

104

105

106

107

In recent years, the task of generating samples from a distribution that characterizes a specific dataset or target image has emerged as a critical challenge in machine learning. This problem has been extensively studied, with solutions primarily leveraging neural networks within deep learning frameworks. Many contemporary generative models operate under the implicit assumption that datasets comprise low-dimensional objects embedded within a high-dimensional space. However, the underlying geometry of the dataset is not always explicitly considered. For instance, variational autoencoders (VAEs) (Kingma & Welling, 2013) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) construct a functional mapping from the low-dimensional latent space to the high-dimensional pixel space. This functional mapping can be interpreted as a transformation from the manifold coordinate system to the pixel coordinate system. More recent approaches, such as diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020a), adopt a more implicit perspective on manifold structure. Geometrically, these models can be viewed as learning a directional field that guides noisy points back to the data manifold, enabling iterative projection. The diffusion process gradually transforms random noise into realistic samples by iteratively denoising along paths that converge onto the data manifold.

A central concept in many of these generative approaches is the Manifold Hypothesis (Loaiza-Ganem et al., 2024), which posits that real-world high-dimensional data, such as images, often concentrates near a low-dimensional manifold embedded within the ambient space. This geometric perspective provides a powerful conceptual framework for understanding generative models and has significantly influenced the design of numerous architectures and training objectives. Various other manifold-aware generative approaches have been proposed. Some methods explicitly model data as residing on specific manifolds. For instance, hyperspherical VAEs (Davidson et al., 2022) and hyperbolic VAEs (Mathieu et al., 2019) adapt generative models to handle data that naturally lies on non-Euclidean manifolds. Riemannian flow models (Gemici et al., 2016; Mathieu & Nickel, 2020) incorporate Riemannian metrics into flow-based models to explicitly account for the intrinsic geometry of the

data manifold. The relationship between manifold structure and probabilistic frameworks remains an active area of research. Normalizing flows (Rezende & Mohamed, 2015) can be interpreted as learning diffeomorphisms between the data manifold and a simple base distribution. Score-based generative models (Song & Ermon, 2020) utilize the score function (the gradient of the log-density) to characterize the data distribution, establishing a direct connection to the geometry of the data manifold. Recent works on denoising diffusion models (Ho et al., 2020b) can also be interpreted as learning a vector field that guides noisy samples back to the data manifold. Despite these advancements, there remains a gap in unifying the geometric and probabilistic perspectives in generative modeling.

This work addresses this gap by providing a geometric interpretation of autoencoders, leveraging geometric properties of the data, specifically the distance function to the manifold. We propose a new generative model that synthesizes both geometric and probabilistic approaches, leading to improved performance in generating high-quality samples. Our approach is based on the premise that the data manifold can be represented as a low-dimensional submanifold embedded within a high-dimensional space. We simultaneously learn both the distance function to this manifold and the probability distribution on it.

# 2 Background and Theoretical Framework

Many generative networks assume that images lie on a lower-dimensional manifold defined according to the latent space representation, which is embedded within a higher-dimensional representation space, such as the pixel space or ambient space. This manifold is explicitly modeled by the decoder in autoencoders (AEs) and variational autoencoders (VAEs), and by the generator in various Generative Adversarial Network (GAN) architectures. In all of these models, the manifold  $\mathcal M$  is represented as a Monge patch. Let the latent space be d-dimensional, parameterized by z, and the pixel space be D-dimensional, parameterized by x, that is (see Fig. 1):

$$G(z) = (x_1(z_1, \dots, z_d), \dots, x_D(z_1, \dots, z_d)).$$

In simple terms, the value at each pixel in the image (or in similar manifold-structured data) is a function of the d parameters z. Another (implicit) way to describe a manifold is as the zero level set of a function. The distance function to the manifold in the ambient (representation) space is well suited for this purpose and is defined as follows:

$$\mathcal{D}_{\mathcal{M}}(x) = \min_{y \in \mathcal{M}} \|x - y\|,\tag{1}$$

where  $\|\cdot\|$  denotes the Euclidean norm. In this high-dimensional representation space, the distance function provides a natural measure of the proximity of a point to the manifold. It is well known that  $\mathcal{D}_{\mathcal{M}}$  satisfies the Eikonal equation  $||\nabla \mathcal{D}_{\mathcal{M}}(x)|| = 1$ , with the natural boundary condition  $\mathcal{D}_{\mathcal{M}}(x) = 0$  for all  $x \in \mathcal{M}$ . Moreover, it is clear that  $-\nabla \mathcal{D}_{\mathcal{M}}(x)$  defines a vector field pointing in the direction of the shortest path to the manifold.

Building on this purely geometric consideration, we introduce a probabilistic model and demonstrate how a deterministic, stepwise diffusion-like model for projection onto the manifold is constructed using the score of the introduced probability function. Let us first assume that the probability of an arbitrary point in the ambient space being a clean image decreases exponentially with the distance from the manifold of clean and meaningful images. We adopt a simple model in which each image x is assumed to be generated from the closest point  $x^*$  on the manifold, with additive Gaussian noise,  $x = x^* + \epsilon$ , such that  $\epsilon \sim \mathcal{N}(0, \sigma_d^2)$ . Therefore, the probability of x to be on the manifold is given by

$$P_d(x) = \frac{1}{Q_d} \exp\left(-\frac{\mathcal{D}_{\mathcal{M}}^2(x)}{2\sigma_d^2}\right),\tag{2}$$

where  $Q_d$  is a normalization factor.

An alternative way to construct a probability distribution based on the manifold hypothesis is to describe each point x in the ambient space as a noisy version of a point (or points) on

the manifold, such that  $x = G(z) + \epsilon$  is conditioned on z. The conditional probability is then given by

$$P(x|G(z)) = \frac{1}{Q_d} \exp\left(-\frac{\|x - G(z)\|^2}{2\sigma_d^2}\right).$$
 (3)

Based on the definition of the distance function, we derive the following expression using the maximum likelihood principle:

$$P_d(x) = \max_z P(x|G(z)) = \frac{1}{Q_d} \exp\left(-\min_z \frac{\|x - G(z)\|^2}{2\sigma_d^2}\right) = \frac{1}{Q_d} \exp\left(-\frac{\mathcal{D}_{\mathcal{M}}^2(x)}{2\sigma_d^2}\right). \tag{4}$$

Note that  $P_d$  implicitly assumes a uniform distribution of data points on the manifold, so the only factor that influences the (conditional) probability is the distance from the image x in the ambient space to the manifold. To account for a non-uniform probability distribution on the manifold, we define the probability function over the ambient space as a marginal distribution (see Appendix A.2 for details):

$$P_{\text{non-u}}(x) = \int_{\mathbb{R}^d} P(x|G(z))P(z)dz,$$

where the subscript "non-u" stands for non-uniform. In this formulation, the probability at x is obtained by integrating contributions from all points on the manifold, where the conditional probability depends solely on the distance to the manifold and is thus purely geometric. Each contribution is weighted by P(z), which represents the likelihood that the point G(z) on the manifold corresponds to a clean image. Since the distribution P(z) is unknown, we estimate it using a kernel density method. Specifically, we define:

$$P(z) \approx P_{\text{ker}}(z) = \frac{1}{Q_{\text{ker}}} \sum_{\alpha \in S} \exp\left(-\frac{||z - z_{\alpha}||^2}{2\sigma_{\text{ker}}^2}\right),$$
 (5)

where S is the set of latent code indices corresponding to clean images, and  $Q_{\rm ker}$  is the normalization constant. Note that  $\sigma_{\rm ker}$  is a hyperparameter that should be chosen carefully. In Fig. 6, we illustrate  $P_{\rm ker}(z)$ . Clearly, the encoding of a generic image x in the latent space, i.e., F(x), may lie in a region with low probability. The probability of a point x being an image depends on its distance to every point on the manifold, weighted by the probability of that point in the latent space. Using this kernel approximation together with the conditional probability from Eq. (3), we can thus approximate the probability function  $P_{\rm non-u}(x)$  as

$$P_{\text{non-u}}(x) \approx \hat{P}_{\text{non-u}}(x) = \frac{1}{Q_d Q_{\text{ker}}} \sum_{\alpha \in S} \int_{\mathbb{R}^d} \exp\left(-\frac{\|x - G(z)\|^2}{2\sigma_d^2}\right) \exp\left(-\frac{\|z - z_\alpha\|^2}{2\sigma_{\text{ker}}^2}\right) dz. \quad (6)$$

## 3 Geometric View of Diffusion Models

Since the domain of both the encoder F and the distance  $\mathcal{D}_{\mathcal{M}}$  is the ambient space  $\mathbb{R}^D$ , effectively training mappings that enable the diffusion-like flow from corrupted images back to clean ones on the manifold requires sampling the high-dimensional ambient space, which is an inherently challenging task due to the curse of dimensionality. Following the approach of diffusion models, we generate ambient samples by adding Gaussian noise to the data points. While this sampling strategy does not cover all possible corruptions, it empirically produces useful mappings. Notably, although the models are trained using Gaussian noise, they generalize well to other types of image corruption during testing.

To connect a corrupted image to its clean projection we use the concept of the score. The score is a D-dimensional vector field defined by  $s(x) = \nabla_x \log P(x)$ , which points in the direction of the steepest ascent of the probability density. For the distance-based probability distribution  $P_d(x)$  defined in Eq. (2), we obtain:

$$s_d(x) = \nabla_x \log P_d(x) = \frac{\nabla_x P_d(x)}{P_d(x)} = -\frac{1}{\sigma_d^2} \mathcal{D}_{\mathcal{M}}(x) \nabla_x \mathcal{D}_{\mathcal{M}}(x). \tag{7}$$

Since  $\mathcal{D}_{\mathcal{M}}(x)$  is the distance to the manifold, its normalized gradient  $\nabla_x \mathcal{D}_{\mathcal{M}}(x)/|\nabla_x \mathcal{D}_{\mathcal{M}}(x)|$  is a unit vector pointing to the closest point on the manifold. Therefore, for  $\sigma_d = 1$  we have:

$$x^{\text{shift}} := x + s_d(x) = x - \mathcal{D}_{\mathcal{M}}(x)\nabla_x \mathcal{D}_{\mathcal{M}}(x)/|\nabla_x \mathcal{D}_{\mathcal{M}}(x)| = G(F(x)) = x^*, \tag{8}$$

where  $x^*$  is the point on the manifold closest to x (see Fig. 2).

To incorporate the probability distribution of clean images on the manifold (or equivalently, in the latent space), we interpret the probability in the ambient space as a marginal distribution. This allows the approximation of the score function using a kernel-based method:

$$s_{\text{non-u}}(x) = \nabla_x \log P_{\text{non-u}} \approx \nabla_x \log \hat{P}_{\text{non-u}} =: \hat{s}_{\text{non-u}}(x).$$

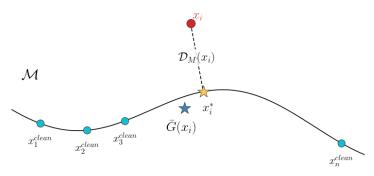


Figure 2: The manifold  $\mathcal{M}$  is illustrated as the curved line.  $x_i^*$  is the closest point to x on the manifold.  $\bar{G}(x)$  is depicted as well and is not necessarily a point on the manifold.

Direct computation results in

$$\hat{s}_{\text{non-u}}(x) = -\frac{1}{2\sigma_d^2} \left( x - \bar{G}(x) \right), \tag{9}$$

where  $\bar{G}(x) = \sum_{\alpha \in S} \bar{G}_{\alpha}(x)$ , and

$$\bar{G}_{\alpha}(x) = \frac{1}{\hat{P}_{\text{non-u}}(x)Q_dQ_{\text{ker}}} \int \left[ G(z)P(x|G(z)) \exp\left(-\frac{||z-z_{\alpha}||^2}{2\sigma_{\text{ker}}^2}\right) \right] dz.$$
 (10)

Note that  $\bar{G}(x)$ , which is the mean of the contributions from all points in the manifold to the probability P(x) does not necessarily lie on the manifold. In contrast,  $x^* = G(F(x))$  is, by definition, a point on the manifold. See Fig. 2 for an illustration and Fig. 8 for a synthetic example. The integral over z in the computation of  $\bar{G}_{\alpha}(x)$  is approximated by randomly sampling the normal distribution centered around the training point  $z_{\alpha}$  (see details in Appendix A).

A noisy or corrupted image x can be viewed as a point in the ambient space. The image generation then becomes the task of finding an appropriate, though not necessarily orthogonal, projection of this point onto the manifold of clean, meaningful images. If the mappings and functions G, F, and  $\mathcal{D}_{\mathcal{M}}$  are perfectly accurate, a single step can move x closer to the corresponding clean image. Since the ambient space is sampled sparsely, especially in regions far from the manifold, the approximations of these mappings become less accurate as the distance from the manifold increases. To address this, we employ multiple iterative steps, gradually improving accuracy as we move closer to the manifold. This process resembles a diffusion-like flow; see Fig. 8 for an illustrative example. Equations 7 and 8 motivate a diffusion-like process guided by the distance function. The score defines a vector field in the ambient space. A step in the direction of the closest point on the manifold is

$$x^{n+1} = x^n - \alpha \mathcal{D}_{\mathcal{M}}(x^n) \nabla_x \mathcal{D}_{\mathcal{M}}(x^n) / |\nabla_x \mathcal{D}_{\mathcal{M}}(x^n)|$$
 with  $0 < \alpha < 1$  and  $x^0 = x$ . (11)

Equation 11 does not take into account the distribution of training points on the manifold. To address this limitation, we combine it with the score of the kernel method to obtain:

$$x^{n+1} = (1 - \beta)x^n + \beta \bar{G}(x^n) - \alpha \mathcal{D}_{\mathcal{M}}(x^n) \nabla_x \mathcal{D}_{\mathcal{M}}(x^n) / |\nabla_x \mathcal{D}_{\mathcal{M}}(x^n)|, \tag{12}$$

where  $0 < \alpha, \beta, \alpha + \beta < 1$ , and  $x^0 = x$ . The trajectory of x as it moves towards the manifold is illustrated in Fig. 8 in Appendix D.

# 4 Methods

## 4.1 Manifold-Probabilistic Projection Model (MPPM)

The autoencoder and the distance function are implemented as separate neural networks and are jointly trained using the loss function in Appendix A.1. Algorithm 1 outlines the training procedure using the clean dataset  $\mathcal{X}^{\text{clean}}$  and the reconstruction of a noisy point x in the ambient space. The algorithm is demonstrated for the simple case of a non-uniform distribution on the circle embedded in  $\mathbb{R}^3$  in Fig. 9 in Appendix D. All the experimental and optimization details can be found in appendices C and D.

## Algorithm 1 MPPM

```
\begin{array}{l} \textbf{function $\operatorname{Train}(\mathcal{X}^{\operatorname{clean}}, \epsilon \sim \mathcal{N}(0, \sigma_d^2))$} \\ G, F, \mathcal{D}_{\mathcal{M}} \leftarrow \operatorname{Train}\big(\mathcal{X}^{\operatorname{clean}}, \epsilon, \mathcal{L}(F, G, \mathcal{D}_{\mathcal{M}})\big)$} \\ \textbf{end function} \\ \textbf{function $\operatorname{RECONSTRUCTION}(x, \mathcal{X}^{\operatorname{clean}}, \alpha, \beta, \operatorname{num\_steps})$} & \rhd 0 < \alpha, \beta, \alpha + \beta < 1$ \\ x^1 \leftarrow x \\ \textbf{for $n \leftarrow 1$ to num\_steps } \textbf{do} \\ x^{n+1} \leftarrow (1-\beta)x^n + \beta \sum_{\alpha} \bar{G}_{\alpha}(x^n) - \alpha \mathcal{D}_{\mathcal{M}}(x^n) \nabla_x \mathcal{D}_{\mathcal{M}}(x^n) / |\nabla_x \mathcal{D}_{\mathcal{M}}(x^n)|$} & \text{by 12, 10} \\ \textbf{end for} \\ \textbf{return $x^{n+1}$} \\ \textbf{end function} \end{array}
```

## 4.2 Latent MPPM (LMPPM)

The key difference between the pixel space and the latent space is that, in the latter, we do not assume that encoded clean and meaningful images lie on a lower-dimensional manifold. Instead, we treat the set of encoded clean and meaningful images as a point cloud that occupies the full dimension of the latent space. We model this set as samples from a probability distribution P(z). Let the set of clean and meaningful images be  $\mathcal{X}^{\text{clean}}$  and the set of these encoded images be  $S = \{F(\mathcal{X}^{\text{clean}})\}$ . In this context, S serves the role that the manifold  $\mathcal{M}$  played in the previous section, in the sense that the distance function  $\mathcal{D}_S$  is now computed in the latent space with respect to the set S. Let  $x \in \mathbb{R}^D$  be an image and  $z = F(x) \in \mathbb{R}^d$  its latent representation. The reconstructed image is then given by  $\hat{x} = G(z)$ . Let us define a distance function  $\mathcal{D}_S : \mathbb{R}^d \to \mathbb{R}$  such that  $\mathcal{D}_S(z)$  measures the distance from z to the set S in the latent space. Using this, we define a shift in the latent space as:  $z^{\text{shift}} := z - \mathcal{D}_S(z) \nabla_z \mathcal{D}_S(z) / |\nabla_z \mathcal{D}_S(z)|$ . The loss function is then given by

$$\mathcal{L}(F, G, \mathcal{D}_{S}) = \lambda_{1} \sum_{z_{i} \notin S} (\mathcal{D}_{S}(z_{i}) - ||z_{i} - z_{i}^{*}||])^{2} + \lambda_{2} \sum_{z_{i} \in S} (x_{i}^{\text{clean}} - G(z_{i}))^{2}$$

$$\lambda_{3} \sum_{z_{i} \in S} |\mathcal{D}_{S}(z_{i})|^{2} + \lambda_{4} \sum_{z_{i}} (\mathcal{D}_{S}(z_{i}) - |\mathcal{D}_{S}(z_{i})|)^{2}$$

$$+ \lambda_{5} \sum_{z_{i} \notin S} ||z_{i}^{\text{shift}} - z_{i}^{*}|| + \lambda_{6} \sum_{z_{i} \notin S} ||G(z_{i}^{\text{shift}}) - x_{i}^{*}||, \tag{13}$$

where  $x_i^* = \arg\min_{\tilde{x} \in \mathcal{X}^{\text{clean}}} \|x_i - \tilde{x}\|$ , and  $z_i^* = F(x_i^*)$ . These definitions ensure that a generic point x in the ambient space, whose closest clean image in the dataset is  $x^*$  is mapped to z = F(x) such that its nearest neighbor in S is  $z^* = F(x^*)$ . It is important to note that the set S evolves over training iterations as the encoder F and decoder F are updated, and the distance function  $\mathcal{D}_S$  is adjusted accordingly. The first three terms are the heart of the algorithm. The 4th element ensures positivity. The 5th and 6th terms improve consistency between all three networks. Ablation study empirically proves that these terms contribute to the performance of the method. By the kernel method, we obtain

$$\bar{z} = \frac{1}{Q} \sum_{x_j \in \mathcal{X}} F(x_j) \exp\left(-\frac{(z - F(x_j))^2}{2\sigma_{ker}^2}\right). \tag{14}$$

The complete procedure is described in Algorithm 2.

# Algorithm 2 LMPPM

```
\begin{array}{l} \textbf{function $\operatorname{Train}(\mathcal{X}^{\operatorname{clean}}, \epsilon \sim \mathcal{N}(0, \sigma_d^2))$} \\ G, F, \mathcal{D}_{\mathcal{S}} \leftarrow \operatorname{Train}\big(\mathcal{X}^{\operatorname{clean}}, \epsilon, \mathcal{L}(F, G, \mathcal{D}_{\mathcal{S}})\big) & \text{by 13} \\ \textbf{end function} \\ \textbf{function $\operatorname{RECONSTRUCTION}(x, \mathcal{X}^{\operatorname{clean}}, \alpha, \beta, \operatorname{num\_steps})$} & \rhd 0 < \alpha, \beta, \alpha + \beta < 1 \\ z^1 \leftarrow F(x) \\ \textbf{for $n \leftarrow 1$ to num\_steps } \textbf{do} \\ z^{n+1} \leftarrow (1-\beta)z^n + \beta\bar{z}^n - \alpha\mathcal{D}_{\mathcal{S}}(z^n)\nabla_z\mathcal{D}_{\mathcal{S}}(z^n)/|\nabla_z\mathcal{D}_{\mathcal{S}}(z^n)|$} & \text{by 14} \\ \textbf{end for} \\ \textbf{return $G^{\hat{c}}(z^{n+1})$} \\ \textbf{end function} \end{array}
```

## 5 Experiments

We evaluated our MPPM method on synthetic data and our LMPPM method on real-world image datasets, where we simultaneously trained an autoencoder-like network for F and G, and a different network for the distance function  $\mathcal{D}_{\mathcal{M}}$  and  $\mathcal{D}_{S}$ . It is important to note that training was performed exclusively with Gaussian noise degradation, while at inference time we evaluated the models under a variety of other degradation types. We compared our results with standard denoising autoencoders (DAE) (Vincent et al., 2008) and latent diffusion models (LDM) (Rombach et al., 2022). For synthetic experiments, we evaluated on a one-dimensional manifold: a half-circle lying in the xy plane and embedded in  $\mathbb{R}^3$ . The points in the circle are sampled according to angular coordinates drawn from truncated normal distributions (see Fig. 8).

For real-world data, we experiment with MNIST (LeCun, 1998) and the SCUT-FBP5500 facial beauty dataset (Liang et al., 2018). To evaluate restoration performance, we apply three types of degradation to MNIST: Gaussian noise, downsampling (super-resolution), and elastic deformation, each at two severity levels. For SCUT-FBP5500, we consider four types of degradation: Gaussian noise, downsampling, random scribbles, and black patches (inpainting), also applied at two severity levels. We train our proposed methods and the comparison baselines to assess their performance across the different datasets. Detailed architecture specifications and hyperparameters are provided in appendices B and C. For synthetic data, we implement MPPM using MLP architectures. For MNIST, we employ a CNN-based autoencoder for both DAE and our LMPPM method, while for SCUT-FBP5500 we adopt a U-Net architecture with skip connections. In addition, we construct an extra set of skip connections from the latent space and combine them with the original skips through weighted summation (see Appendix B). The distance functions  $\mathcal{D}_{\mathcal{M}}$  and  $\mathcal{D}_{\mathcal{S}}$  are implemented as MLPs with progressively decreasing layer sizes to perform dimensionality reduction. For LDM, we integrate the corresponding DAE backbone (in place of the autoencoder) with a standard diffusion model, using 2000 diffusion steps.

## 5.1 Results

**MNIST Results:** For the MNIST dataset, we set the latent space dimension to 18 and the additive noise to  $\epsilon = 0.4$ . To calculate FID, we trained an MNIST classifier and computed an

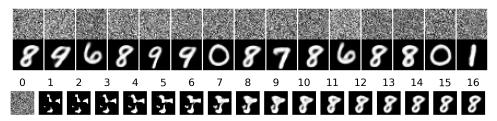


Figure 3: Top: Digit generation from pure noise, with an FID of 19.53 computed over 2000 images. Bottom: Progression of digit generation over 16 steps.

Table 1: Quantitative results on MNIST

	$ \begin{array}{ c c c } \hline Elastic 2.3 \\ SSIM \uparrow & FID \downarrow \end{array} $		Elastic 1.8 SSIM $\uparrow$ FID $\downarrow$	
DAE LDM LMPPM (ours)	$\begin{array}{c} 0.66 \\ 0.64 \\ 0.63 \end{array}$	69.36 66.52 <b>12.61</b>	$\begin{array}{c c} 0.59 \\ 0.58 \\ 0.59 \end{array}$	134.60 124.05 <b>16.38</b>
	Downsan SSIM ↑	nple 0.5 FID↓	Downsan SSIM ↑	$\begin{array}{c} \text{ple } 0.35 \\ \text{FID } \downarrow \end{array}$
DAE LDM LMPPM (ours)	$\begin{array}{c} 0.79 \\ 0.75 \\ 0.67 \end{array}$	31.66 31.61 <b>11.27</b>	$\begin{array}{ c c }\hline 0.54 \\ 0.53 \\ 0.52 \\ \end{array}$	133.66 128.80 <b>22.65</b>

Table 2: Quantitative results on SCUT-FBP5500

	Miss pix   SSIM ↑	els $0.04$ FID $\downarrow$	Miss pix SSIM ↑	els $0.08$ FID $\downarrow$	Miss pixe	els 0.1 FID↓
DAE LDM LMPPM (ours)	$\begin{array}{ c c } \hline 0.917 \\ \hline 0.914 \\ 0.881 \\ \hline \end{array}$	33.90 27.35 <b>16.20</b>	$\begin{array}{c} 0.798 \\ 0.798 \\ \underline{0.862} \end{array}$	49.00 41.47 <b>23.92</b>	$ \begin{array}{c c} 0.745 \\ 0.738 \\ 0.832 \end{array} $	47.94 44.41 <b>34.13</b>
	Scribl	ole 6 FID↓	Scribb SSIM ↑	ole 13 FID↓	Scribbl SSIM ↑	e 20 FID↓
DAE LDM LMPPM (ours)	$\begin{array}{ c c }\hline 0.921\\ 0.919\\ 0.879\\ \end{array}$	34.83 29.31 <b>16.73</b>	$\begin{array}{c} 0.889 \\ 0.887 \\ 0.878 \end{array}$	45.66 39.02 <b>17.35</b>	0.860 0.859 0.869	51.68 44.66 <b>18.46</b>
	Sharp SSIM ↑	en 8 FID↓	Sharpe SSIM ↑	en 10 FID↓	Sharpe SSIM ↑	n 18 FID↓
DAE LDM LMPPM (ours)	$\begin{array}{ c c }\hline 0.902\\ 0.898\\ 0.878\\ \end{array}$	28.53 20.79 <b>16.79</b>	$\begin{array}{c c} 0.883 \\ 0.878 \\ 0.874 \end{array}$	29.80 21.73 <b>17.33</b>	0.815 0.807 0.853	33.82 25.37 <b>19.48</b>

embedding distribution for each class. After reconstructing a degraded digit, we classified it and compared its embedding with the corresponding pre-computed class distribution. Table 1 reports the mean SSIM and FID metrics. Our method consistently outperforms both DAE and LDM baselines across all degradation types in terms of FID scores. Notably, DAE occasionally achieved higher SSIM values, although its visual results were inferior. Fig. 10 in Appendix D illustrates restoration examples for Gaussian noise, elastic deformation, and downsampling. Additional experiment included the generation of digits from a pure noise. We generated 200 images from random Gaussian noise and managed to obtain realistic digits (FID=19.5) as can be seen in Fig. 3.

SCUT-FBP5500 Results: Figure 4 shows restoration results on facial images with several degradation functions: excessive Gaussian noise, randomly missing pixels, random scribbles and over sharpening. The quantitative results in Table 2 support these visual observations, with our approach achieving consistently lower FID values across all degradation types. We set the latent dimension to 1024 and the additive noise to  $\epsilon = 0.2$ . While in some cases the DAE method achieves higher SSIM values, the visual quality of its reconstructions is noticeably inferior. Additional results for Gaussian noise, downsampling, and over-sharpening are provided in Appendix D.

#### Summary and Conclusions

This work emphasizes the Manifold Hypothesis and interprets established image restoration and generation methods through a novel geometric perspective. Beyond presenting a

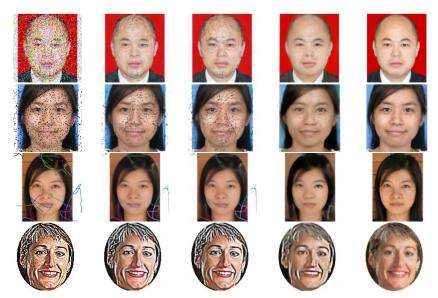


Figure 4: Variety of degradations: noise, missing pixels, scribbles and over sharpening. Left to right: degraded, DAE, LDM, LMPPM (ours), and original.



Figure 5: Left to right: over-sharpened input, DAE, LDM, LMPPM (ours), and original. LMPPM remains realistic despite changes to the face.

unifying framework, which is valuable in its own right, we propose incorporating a learned distance function to the manifold. By leveraging distances to the manifold, we establish a connection between the geometric structure and a probability density approximation. By employing a kernel-like method to approximate the probability distribution on the manifold, or equivalently on the latent space, we integrate geometry and probability in a novel manner. We induce a vector field in the ambient space via the score of these probability densities. This vector field directs each point toward the manifold of clean images, considering both the structure and the distribution of clean and meaningful images on the manifold.

In this work, we utilize a (denoising) autoencoder in conjunction with the distance function. Providing an approach where both F and G define the manifold while maintaining their coupling to the distance function  $\mathcal{D}$  from it. However, due to potential errors in the outputs of the three networks G, F and  $\mathcal{D}$ , especially when x is far from the manifold, this vector field is not exact. Therefore, rather than applying a single-step (weighted) projection onto the manifold, we proceed iteratively, advancing in small steps along the noisy vector field. We are currently exploring an analogous approach where VAE and GAN are coupled with the distance function. A key practical advantage of our approach is its application in the latent space. This dimensionality reduction significantly enhances the accuracy of the distance function, thereby improving restoration and generation results. Indeed, as shown in our experiments (Section 5.1), comparisons with other leading methods indicate the superior performance of our methods, particularly under severe distortions for different data sets and different distortions.

## References

- Tim R. Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M. Tomczak. Hyperspherical variational auto-encoders, 2022. URL https://arxiv.org/abs/1804.00891.
  - Mevlana C Gemici, Danilo Jimenez Rezende, and Shakir Mohamed. Normalizing flows on Riemannian manifolds. In *NeurIPS Workshop on Bayesian Deep Learning*, 2016.
    - Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
  - Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Advances in Neural Information Processing Systems, 2020a.
  - Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020b.
  - Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2013. URL https://arxiv.org/abs/1312.6114.
  - Yann LeCun. The mnist database of handwritten digits. http://yann. lecun. com/exdb/mnist/, 1998.
  - Lingyu Liang, Luojun Lin, Lianwen Jin, Duorui Xie, and Mengru Li. Scut-fbp5500: A diverse benchmark dataset for multi-paradigm facial beauty prediction. 2018.
  - Gabriel Loaiza-Ganem, Brendan Leigh Ross, Rasa Hosseinzadeh, Anthony L. Caterini, and Jesse C. Cresswell. Deep generative models through the lens of the manifold hypothesis: A survey and new connections, 2024. URL https://arxiv.org/abs/2404.02954.
  - Emile Mathieu and Maximilian Nickel. Riemannian continuous normalizing flows. In Advances in Neural Information Processing Systems, 2020.
  - Emile Mathieu, Charline Le Lan, Chris J. Maddison, Ryota Tomioka, and Yee Whye Teh. Continuous hierarchical representations with poincaré variational auto-encoders, 2019. URL https://arxiv.org/abs/1901.06033.
  - Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, 2015.
  - Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
  - David E. Rumelhart and James L. McClelland. Learning Internal Representations by Error Propagation, pp. 318–362. 1987.
- Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics, 2015. URL https://arxiv.org/abs/1503.03585.
  - Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution, 2020. URL https://arxiv.org/abs/1907.05600.
  - Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103, 2008.

A APPENDIX: DETAILED THEORY

#### A.1 MPPM

The loss function for the MPPM is

$$\mathcal{L}(F, G, \mathcal{D}_{\mathcal{M}}) = \lambda_{1} \sum_{x_{i} \notin \mathcal{M}} (\mathcal{D}_{\mathcal{M}}(x_{i}) - \|x_{i} - x_{i}^{*}\|])^{2} + \lambda_{2} \sum_{x_{i} \in \mathcal{M}} \left(x_{i}^{\text{clean}} - G(F(x_{i}^{\text{clean}}))^{2} + \lambda_{3} \sum_{x_{i} \in \mathcal{M}} |\mathcal{D}_{\mathcal{M}}(x_{i})|^{2} + \lambda_{4} \sum_{x_{i} \in \mathbb{R}^{D}} (\mathcal{D}_{\mathcal{M}}(x_{i}) - |\mathcal{D}_{\mathcal{M}}(x_{i})|)^{2} + \lambda_{5} \sum_{x_{i} \in \mathbb{R}^{D}} \left(x_{i}^{\text{shift}} - x_{i}^{*}\right)^{2},$$

$$(15)$$

where  $x^* = G(F(x))$ . The first term defines the distance function assuming a perfect autoencoder; the second is the standard autoencoder loss. The third term enforces the boundary condition on the distance function and the fourth ensures its positivity. The last term enforces the geometric consistency of Eq. (8) (see also Fig. 1).

#### A.2 Kernel method

Detailed computation of eq. 2

$$\begin{split} P_{\text{non-u}}(x) &= \int_{\mathcal{M}} P(x,y=G(z)) dy = \int_{\mathcal{M}} P(x|y=G(z)) P(y) dy \\ &= \int_{\mathbb{R}^d} P(x|G(z)) P(G(z)) \sqrt{g} dz = \int_{\mathbb{R}^d} P(x|G(z)) P(z) dz, \end{split}$$

where the subscript "non-u" stands for non-uniform. Here  $dy = \sqrt{g}dz$  is the manifolds volume element, where  $g = \det G$  and  $G_{\mu\nu} = \sum_{i=1}^{D} J^i_{\mu}J^i_{\nu}$  is the induced metric, with the Jacobian of the embedding map given by  $J^i_{\mu} = \partial G^i(z)/\partial z_{\mu}$ . In the last equality, we use the identity  $P(G(z)) = P(z)(\sqrt{g})^{-1}$ .

#### A.3 SCORE FUNCTION

$$s_{\text{non-u}}(x) = \nabla_x \log P_{\text{non-u}} \approx \nabla_x \log \hat{P}_{\text{non-u}} =: \hat{s}_{\text{non-u}}(x).$$

Direct computation results in

$$\hat{s}_{\text{non-u}}(x) = -\frac{1}{2\sigma_{J}^{2}} \left( x - \bar{G}(x) \right), \tag{16}$$

where  $\bar{G}(x) = \sum_{\alpha \in S} \bar{G}_{\alpha}(x)$ , and

$$\bar{G}_{\alpha}(x) = \frac{1}{\hat{P}_{\text{non-u}}(x)Q_dQ_{\text{ker}}} \int \left[ G(z)P(x|G(z)) \exp\left(-\frac{||z - z_{\alpha}||^2}{2\sigma_{\text{ker}}^2}\right) \right] dz.$$
 (17)

Specifically,

$$\hat{s}_{\text{non-u}}(x) = \frac{1}{\hat{P}_{\text{non-u}}} \nabla_x \hat{P}_{\text{non-u}} = \frac{1}{\hat{P}_{\text{non-u}}} \nabla_x \left( \int P(x \mid G(z)) P_{\text{ker}}(z) dz \right).$$

Now,

$$(\nabla_x P(x \mid G(z))) P_{\text{ker}}(z) = -\frac{1}{2\sigma_d^2 Q_d} (x - G(z)) \exp\left(-\frac{\|x - G(z)\|^2}{2\sigma_d^2}\right) \frac{1}{Q_{\text{ker}}} \sum_{\alpha \in S} \exp\left(-\frac{\|z - z_\alpha\|^2}{2\sigma_{ker}^2}\right).$$
(18)

The integral of z in the computation of  $\bar{G}_{\alpha}(x)$  is approximated by randomly sampling the normal distribution centered around the training point  $z_{\alpha}$ . Explicitly, we approximate the mean using and average over n samples from  $P_{\ker}$ 

$$\int \left[ G(z)P(x|G(z)) \exp\left(-\frac{||z-z_{\alpha}||^2}{2\sigma_{\ker}^2}\right) \right] dz \approx \frac{1}{n} \sum_{z_i \in \mathcal{N}(z_{\alpha}, \sigma_{\ker}^2)} G(z_i) \exp\left(-\frac{||x-G(z_i)||^2}{2\sigma_d^2}\right), \tag{19}$$

where  $\alpha$  denotes an index in the training set (see Fig. 6). The calculation of  $\bar{G}_{\alpha}(x)$  requires evaluating  $P_{\text{non-u}}(x)$  in the denominator. In particular, we approximate

$$\int \left[ P(x|G(z)) \exp\left(-\frac{||z - z_{\alpha}||^2}{2\sigma_{\ker}^2}\right) \right] dz \approx \frac{1}{n} \sum_{z_i \in \mathcal{N}(z_{\alpha}, \sigma_{\ker}^2)} \exp\left(-\frac{||x - G(z_i)||^2}{2\sigma_d^2}\right). \tag{20}$$

Note that in the computation of  $\bar{G}_{\alpha}$  all constant factors  $Q_d$ ,  $Q_{\text{ker}}$  and  $\frac{1}{n}$ , are canceled between the numerator and the denominator.

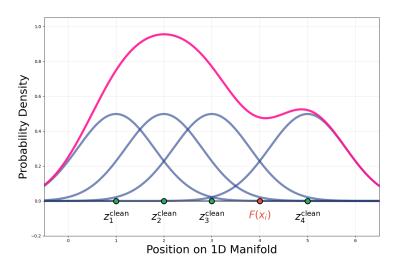


Figure 6: An illustration of the kernel approximation  $P_{\text{ker}}(z)$  of the probability distribution P(z) in the latent space.

## B APPENDIX: DETAILED EXPERIMENTAL SETUP

#### B.1 Notation and Abbreviations

Table 3: Glossary of abbreviations and terms used throughout the paper

Term	Definition
DAE	Denoising Autoencoder
MPPM	Manifold Projection and Propagation Method (our proposed approach)
LMPPM	Latent Manifold Projection and Propagation Method (our proposed approach)
LDM	Latent Diffusion Model
SSIM	Structural Similarity Index Measure
BN	Batch Normalization

Table 4: Summary of experimental datasets used for evaluating restoration performance

Dataset	Description
MNIST	$60,000 \text{ training}/10,000 \text{ test grayscale images } (28 \times 28 \text{ pixels})$
SCUT-FBP5500	$5,500$ facial images with beauty scores (resized to $120 \times 120$ )

#### B.2 Degradations

**Degradation Parameters** We apply six degradation types to simulate real-world image corruption scenarios. Each degradation is applied at three severity levels (mild, intermediate, and severe) to test the robustness of restoration methods:

Table 5: Degradation parameters at different severity levels. Note: Lower  $\sigma$  values for elastic deformation indicate more severe distortion due to increased localized displacement

Degradation	Mild	Severe
Gaussian noise $(\sigma)$	0.2	0.3
Elastic deformation $(\sigma)$	1.5	1.1
Super-resolution (downsampling factor)	0.5	0.35
Missing Pixels	0.04	0.1
Number of Scibbles	13	20
Over Sharpening	10	18

**Degradation Methods** Brief descriptions of each degradation type:

- Gaussian noise: Additive zero-mean Gaussian noise that simulates sensor noise or transmission errors.
- Elastic deformation: Non-rigid distortions implemented using torchvision.transform.ElasticTransform( $\alpha=34,\sigma$ ) that simulate warping effects
- Super-resolution: Downsampling followed by upsampling to original resolution, simulating reconstruction from low-resolution data.
- Missing Pixels: Set black patches with some coverage portion;
- Scribbles Add n random scribbles with random colors
- Over Sharpening by factor  $s: I = I + s(I I * \sigma_s)$

#### B.3 Model Architectures

We implemented three main architectures across all experiments, with design choices tailored to each dataset's complexity.

**Synthetic Data Model** Synthetic data for MPPM experiments use MLP-based networks with a latent dimension of 8, selected based on the low intrinsic dimensionality of these manifolds:

Table 6: Network architectures for synthetic data experiments. All models use fully-connected layers

Component	Architecture
Encoder	$3 \rightarrow 64 \rightarrow 32 \rightarrow 16 \rightarrow 8$ with ReLU
Decoder	$8 \rightarrow 16 \rightarrow 32 \rightarrow 64 \rightarrow 3$ with ReLU
Distance Network	$8 \rightarrow 64 \rightarrow 32 \rightarrow 16 \rightarrow 1$ with ReLU, dropout=0.2

MNIST Models MNIST experiments use CNN-based models with latent dimension 18, chosen to capture the variability among handwritten digits while promoting compact representations.

Table 7: Network architectures for MNIST experiments

Component	Architecture
Encoder	$\operatorname{Conv2d}(1 \to 32 \to 64, \operatorname{kernel} = 3, \operatorname{stride} = 2) \to \operatorname{Flatten} \to \operatorname{Linear}(64 \times 7 \times 7 \to 18)$
Decoder	$Linear(18 \rightarrow 64 \times 7 \times 7) \rightarrow Reshape \rightarrow ConvTranspose2d(64 \rightarrow 32 \rightarrow 1) \rightarrow Sigmoid$
Distance Network	$18 \rightarrow 100 \rightarrow 50 \rightarrow 20 \rightarrow 1$ with ReLU, dropout=0.2

**SCUT-FBP5500 Models** Facial image experiments employ a U-Net with skip connections and a latent dimension of 1024, which accommodates the higher complexity of facial features while enabling detailed reconstruction. Facial image experiments employ a U-Net

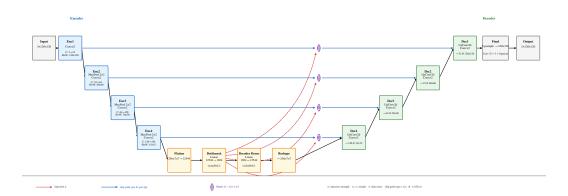


Figure 7: Modified U-net architecture

with skip connections and a latent dimension of 1024, which accommodates the higher complexity of facial features while enabling detailed reconstruction. Note that in the U-Net architecture, during the inference process we use iterations (denoted by superscripts) such that

 $F(x^n) = (S_1^n, S_2^n, \dots, S_k^n, z^n)^T,$ 

and

$$x^{n+1} = G\left(S_1^n + \hat{S}_1^n(z^{n+1}), S_2^n + \hat{S}_2^n(z^{n+1}), \dots, S_k^n + \hat{S}_k^n(z^{n+1}), z^{n+1}\right).$$

Here, each  $\hat{S}^n_i(z^{n+1})$  denotes the projection of the latent space  $z^{n+1}$  onto the corresponding skip connection  $S^n_i$ . Thus, the updated skip connection is formed by adding the original skip feature  $S^n_i$  with the new projected feature  $\hat{S}^n_i(z^{n+1})$  before being passed to G. The architecture is illustrated in Fig. 7.

Component	Structure
EncoderBlock $(C_{in} \to C_{out})$	$\operatorname{Conv}(C_{in} \to C_{out}) \to \operatorname{BN} \to \operatorname{LReLU} \to$
	$\operatorname{Conv}(C_{out} \to C_{out}) \to \operatorname{BN} \to \operatorname{LReLU} \to \operatorname{MPool}$
Encoder	EncoderBlock(3 $\rightarrow$ 32), output: $60 \times 60$
	EncoderBlock(32 $\rightarrow$ 64), output: 30 $\times$ 30
	EncoderBlock(64 $\rightarrow$ 128), output: 15 $\times$ 15
	EncoderBlock(128 $\rightarrow$ 256), output: 7 $\times$ 7
	$Flatten \rightarrow Linear(12544 \rightarrow 1024) \rightarrow LReLU$
$- \text{DecoderBlock}(C_{in}, C_{skip}, C_{out})$	$\operatorname{ConvT}(C_{in} \to C_{in}) \to \operatorname{Cat}([C_{in}, C_{skip}]) \to$
	$\operatorname{Conv}(C_{in} + C_{skip} \to C_{in}) \to \operatorname{BN} \to \operatorname{LReLU} \to$
	$\operatorname{Conv}(C_{in} \to C_{out}) \to \operatorname{BN} \to \operatorname{LReLU}$
Decoder	$\operatorname{Linear}(1024 \to 12544) \to \operatorname{Reshape}(256, 7, 7)$
	DecoderBlock(256, 256, 128), output: $15 \times 15$
	DecoderBlock(128, 128, 64), output: $30 \times 30$
	DecoderBlock(64, 64, 32), output: $60 \times 60$
	DecoderBlock $(32, 32, 32)$ , output: $120 \times 120$
	$Conv(32 \to 3) \to Sigmoid$
Distance Network	$1024 \rightarrow 100 \rightarrow 50 \rightarrow 20 \rightarrow 1$ with ReLU, dropout=0.2

Table 8: Network architectures for SCUT-FBP5500 experiments. Skip connections connect corresponding Encoder and Decoder layers through concatenation. The encoder and decoder blocks are represented as parameterized functions (shown in italic font), where  $C_{in}$ ,  $C_{out}$ , and  $C_{skip}$  represent the number of input, output, and skip connection channels respectively. Abbreviations: Conv = Conv2d (kernel=3, padding=1), BN = BatchNorm2d, LReLU = LeakyReLU(0.2), MPool = MaxPool2d(2), ConvT = ConvTranspose2d(kernel=2, stride=2), Cat = Concatenation. The bottleneck dimension is 1024.

# C TRAINING AND EVALUATION

Table 9: MPPM training and inference parameters for synthetic data

Parameter	Value
Optimizer	Adam $(\beta_1 = 0.9, \beta_2 = 0.999)$
Learning rates	AE: $1 \times 10^{-3}$ , Distance network: $1 \times 10^{-3}$
Weight decay	$1 \times 10^{-4}$
Batch size	550
Training epochs	500
Loss function	Composite loss (Equation 15)
Early stopping	Patience: 100 epochs
$\alpha$ (distance gradient step)	0.15
$\beta$ (kernel averaging weight)	0.1
Convergence tolerance	0.005
Maximum iterations	60

Table 10: LMPPM training parameters across all experiments, determined through preliminary grid search, diffusion steps are define the number of steps in algorithm 2

Parameter	Value
Optimizer	Adam $(\beta_1 = 0.9, \beta_2 = 0.999)$
Learning rates	AE: $1 \times 10^{-3}$ , Distance network: $1 \times 10^{-5}$ , LDM: $1 \times 10^{-3}$
Batch size	MNIST: 128, SCUT-FBP5500: 32
Training epochs	MNIST: 100, SCUT-FBP5500: 75
Loss functions	DAE: L2, LDM: MSE, LMPPM: Composite loss 13
Early stopping	Patience: 8 epochs
Diffusion steps	MNIST: 2000, SCUT-FBP5500: 2000

## D Additional Results

 Here, we present additional experimental results. Figures 8 and 9 depict the results of the MPPM algorithm.

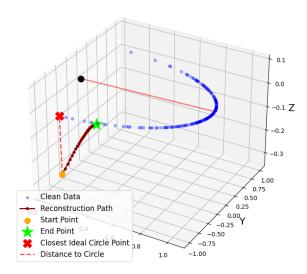
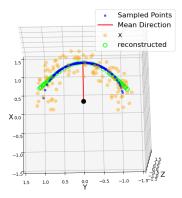
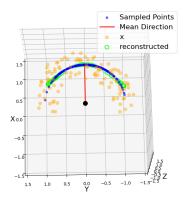


Figure 8: The manifold  $\mathcal{M}$  is the unit circle lying in the xy-plane and is parametrized by the azimuth angle  $\theta$ . It is sampled according to a normal distribution centered at  $\theta_0$  indicated by the red line. The reconstruction trajectory is shown in dark red. Note that the final result of the iterations on x does not converge to  $x^*$  which is the closest point on the circle. Instead, it is influenced by the data distribution on the manifold through the effect of  $\bar{G}(x)$ .

Next, we present additional results on the MNIST dataset using the LMPPM algorithm. Figure 10 shows reconstruction results under noise, elastic, and downsampling deformations, compared with the DAE and LDAM models. Figures 11, 12, 13, and 14 present reconstruction results for missing pixels, scribbles, noise, and over-sharpening deformations, respectively for the SCUT-FBP5500 dataset. We compare our method with the DAE and LDM models. Finally, Figure 15 shows the reconstruction after 4 iterations.





- (a) DAE restoration. MSE = 0.032, max error = 0.147.
- (b) MPPM restoration. MSE = 0.026, max error = 0.060.

Figure 9: Comparison between the DAE and our proposed MPPM, this example uses the same setup as in Fig. 8. The error was computed as the deviation from the unit circle in 2D. In regions of the circle with lower probability density, the DAE is more prone to error than the proposed MPPM method.

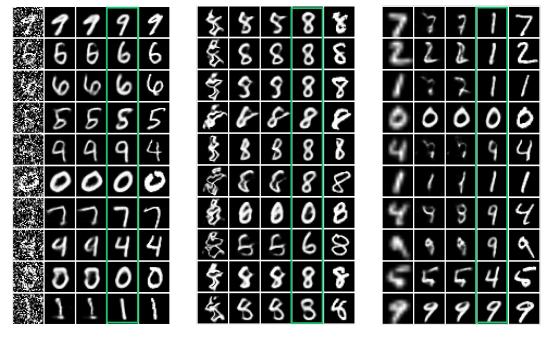


Figure 10: Left panel: noise = 0.7; middle panel: elastic ( $\alpha$  = 0.34,  $\sigma$  = 1.8); right panel: downsampling factor = 0.35. In all panels, from left to right: degraded, DAE, LDA, LMMPM (ours), and original.

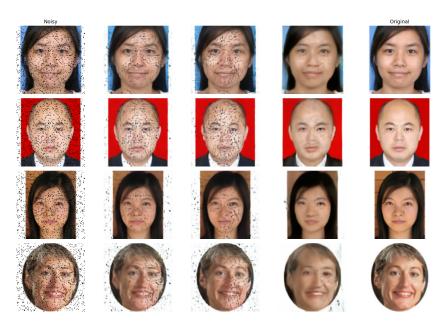


Figure 11: Missing pixels. From left to right: degraded, DAE, Diffusion, LMPPM, original.

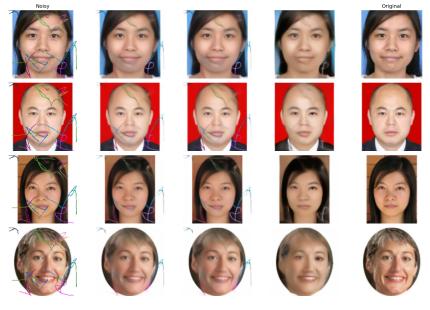


Figure 12: 13 scribbles. From left to right: degraded, DAE, Diffusion, LMPPM, original.

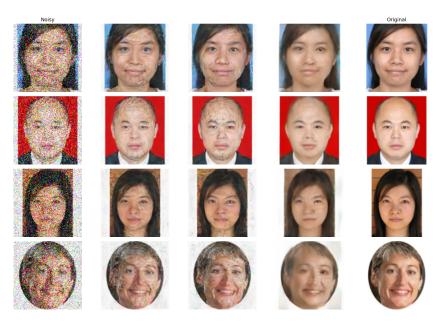


Figure 13: Noise  $\sigma = 0.3$ . From left to right: degraded, DAE, Diffusion, LMPPM, original.

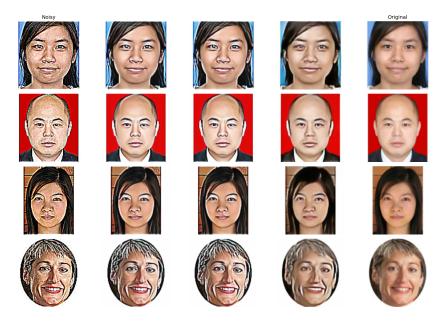


Figure 14: over sharpening. From left to right: degraded, DAE, Diffusion, LMPPM, original.



Figure 15: gradual reconstruction of missing pixels degradation.

## E USE OF LARGE LANGUAGE MODELS

Large Language Models (LLMs) were used in this work solely as a language assistance tool for English polishing and proofreading. Specifically, we employed LLMs to:

- Improve grammar, syntax, and sentence structure in the manuscript
- Enhance clarity and readability of technical descriptions
- Correct spelling and typographical errors
- Suggest more precise word choices and phrasing

The LLMs did not contribute to research ideation, methodology development, experimental design, data analysis, or the generation of scientific content. All research concepts, approaches, results, and conclusions presented in this paper are entirely the work of the human authors. The LLMs were used exclusively for language refinement of content that was already conceptualized and written by the authors.

Additionally, an LLM was used to assist in drafting this disclosure section itself, based on the authors' description of how LLMs were employed in the research process.

We take full responsibility for all content in this manuscript, including any text that was refined with LLM assistance. All factual claims, scientific interpretations, and conclusions remain our own work and responsibility.