
Autoencoder Image Interpolation by Shaping the Latent Space

Alon Oring¹ Zohar Yakhini¹ Yacov Hel-Or¹

Abstract

Autoencoders represent an effective approach for computing the underlying factors characterizing datasets of different types. The latent representation of autoencoders have been studied in the context of enabling interpolation between data points by decoding convex combinations of latent vectors. This interpolation, however, often leads to artifacts or produces unrealistic results during reconstruction. We argue that these incongruities are due to the structure of the latent space and because such naively interpolated latent vectors deviate from the data manifold. In this paper, we propose a regularization technique that shapes the latent representation to follow a manifold that is consistent with the training images and that drives the manifold to be smooth and locally convex. This regularization not only enables faithful interpolation between data points, as we show herein, but can also be used as a general regularization technique to avoid overfitting or to produce new samples for data augmentation.

1. Introduction

Given a set of data points, data interpolation or extrapolation aims at predicting novel data points between given samples (interpolation) or predicting novel data outside the sample range (extrapolation). Faithful data interpolation between sampled data can be seen as a measure of the generalization capacity of a learning system (Berthelot et al., 2018). In the context of computer vision and computer graphics, data interpolation may refer to generating novel views of an object between two given views or predicting in-between animated frames from key frames.

Interpolation that produces novel views of a scene requires input such as the geometric and photometric parameters of

existing objects, camera parameters and additional scene components, such as lighting and the reflective characteristics of nearby objects. Unfortunately, these characteristics are not always available or are difficult to extract in real-world scenarios. Thus, in such cases, we can apply *data-driven interpolation* that is deduced based on a sampled dataset drawn from the scene taken under various acquisition parameters.

The task of data interpolation is to extract new samples (possibly continuous) between known data samples. Clearly, linear interpolation between two images in the input (image) domain does not work as it produces a cross-dissolve effect between the intensities of the two images. Adopting the manifold view of data (Goodfellow et al., 2016; Verma et al., 2018; Bengio et al., 2013), this task can be seen as sampling new data points along the geodesic path between the given points. The problem is that this manifold is unknown in advance and one has to approximate it from the given data. Alternatively, adopting the probabilistic perspective, interpolation can be viewed as drawing samples from highly probable areas in the data space.

One fascinating property of unsupervised learning is the network’s ability to reveal the underlying factors controlling a given dataset. Autoencoders (Doersch, 2016; Kingma & Welling, 2013; Goodfellow et al., 2016; Kramer, 1991; Vincent et al., 2010) represent an effective approach for exposing these factors. Researchers have demonstrated the ability to interpolate between data points by decoding a convex sum of latent vectors (Shu et al., 2018; Mathieu et al., 2016); however, this interpolation often incorporates




Figure 1. Left: A vertical pole casting a shadow. Yellow blocks-top row: Cross-dissolve phenomena as a result of linear interpolation in the input space. Yellow blocks-bottom row: Image reconstruction obtained by a linear latent space interpolation of an autoencoder. Unrealistic artifacts are introduced.

¹School of Computer Science, The Interdisciplinary Center, Herzliya, Israel . Correspondence to: Y.H. <toky@idc.ac.il>.

visible artifacts during reconstruction.

To illustrate the problem, consider the following example: A scene is composed of a vertical pole at the center of a flat plane (Figure 1-left). A single light source illuminates the scene and accordingly, the pole projects a shadow onto the plane. The position of the light source can vary along the upper hemisphere. Hence, the underlying parameters controlling the generated scene are (θ, ϕ) , the elevation and azimuth, respectively. The interaction between the light and the pole produces a cast shadow whose direction and length are determined by the light direction. A set of images of this scene is acquired from a fixed viewing position (from above) with various lighting directions. Our goal in this example is to train a model that is capable of interpolating between two given images. Figure 1, top row, depicts a set of interpolated images, between the source image (left image) and the target image (right image), where the interpolation is performed in the input domain. As illustrated, the interpolation is not natural as it produces cross-dissolve effects in image intensities. Training a standard autoencoder and applying linear interpolation in its latent space generates images that are much more realistic (Figure 1, bottom row). Nevertheless, this interpolation is not perfect as visible artifacts occur in the interpolated images. The source of these artifacts can be investigated by closely inspecting the 2D manifold embedded in the latent space.

Figure 2 shows two manifolds embedded in latent spaces, one with data embedded in 2D latent space (upper left plot) and one with data embedded in 3D latent space (upper right plot). In both cases, the manifolds are 2D and are generated using vanilla autoencoders. The grid lines represent the (θ, ϕ) parameterization. It can be seen that the encoders produce non-smooth and non-convex surfaces in 2D as well as in 3D. Thus, linear interpolation between two data points inevitably produces in-between points outside of the manifold. In practice, the decoded images of such points are unpredictable and may produce non-realistic artifacts. This issue is demonstrated in the two right images in Figure 2. When the interpolated point is on the manifold (an empty circle denoted ‘A’), a faithful image is generated by the decoder (lower left image). When the interpolated point departs from the manifold (the circle denoted ‘B’), the resulting image is unpredictable (lower right image).

In this paper, we argue that the common statistical view of autoencoders is not appropriate when dealing with data that have been generated from continuous factors. Alternatively, the manifold structure of continuous data must be considered, taking into account the geometry and shape of the manifold. Accordingly, we propose a new interpolation regularization mechanism consisting of an *adversarial loss*, a *cycle-consistency loss*, and a *smoothness loss*. The adversarial loss drives the interpolated points to look realistic as it

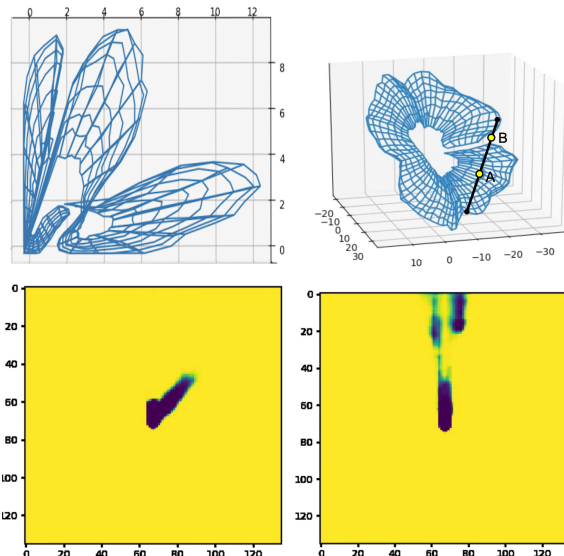


Figure 2. The latent manifold of the data embedded in 2D latent space (upper left plot) and 3D latent space (upper right plot) learned by vanilla autoencoders. Gridlines represent the (θ, ϕ) parameterization. The lower left image was generated from the latent point denoted ‘A’. The lower right image was generated from the latent point denoted ‘B’.

is optimized against a discriminator that learns to tell apart real from interpolated data points. The cycle-consistency and the smoothness losses encourage smooth interpolations between data points. We show empirically that these combined losses prompt the autoencoder to produce realistic and smooth interpolations while providing a convex latent manifold with a bijective mapping between the input and the latent manifolds. This regularization mechanism not only enables faithful interpolation between data points, but can also be used as a general regularization technique to avoid overfitting or to produce new samples for data augmentation, as suggested, among others, by (Zhang et al., 2018).

To conclude, the contributions of the papers are: I. We define what constitutes an admissible interpolation between two data points on a continuous manifold. In particular we added the cycle-consistency and the smoothness terms and show their importance in generating admissible interpolations. II. We empirically demonstrate how the combination of the four losses; the reconstruction, adversarial, cycle-consistency and the smoothness losses, contribute to admissible interpolations and produce state of the art results.

2. Manifold Data Interpolation

Before presenting the proposed approach we would like to define what constitutes a proper interpolation between two data points. There are many possible paths between two points on the manifold. Even if we require the interpolations

to be on a geodesic path, there might be infinitely many such paths between two points. Therefore, we relax the geodesic requirement and define less restrictive conditions. Formally, assume we are given a dataset sampled from a target domain \mathcal{X} . We are interested in interpolating between two data points \mathbf{x}_i and \mathbf{x}_j from \mathcal{X} . Let the interpolated points be $\hat{\mathbf{x}}_{i \rightarrow j}(\alpha)$ for $\alpha \in [0, 1]$ and let $P(\mathbf{x})$ be the probability that a data point \mathbf{x} belongs to \mathcal{X} . We define an interpolation to be an *admissible interpolation* if $\hat{\mathbf{x}}_{i \rightarrow j}(\alpha)$ satisfies the following conditions:

1. **Boundary conditions:**

$$\hat{\mathbf{x}}_{i \rightarrow j}(0) = \mathbf{x}_i \quad \text{and} \quad \hat{\mathbf{x}}_{i \rightarrow j}(1) = \mathbf{x}_j$$

2. **Monotonicity:** We require that under some defined distance on the manifold $d(\mathbf{x}, \mathbf{x}')$, the interpolated points will depart from \mathbf{x}_i and approach \mathbf{x}_j , as the parameterization α goes from 0 to 1. Namely, $\forall \alpha' \geq \alpha$,

$$d(\hat{\mathbf{x}}_{i \rightarrow j}(\alpha), \mathbf{x}_i) \leq d(\hat{\mathbf{x}}_{i \rightarrow j}(\alpha'), \mathbf{x}_i)$$

and similarly:

$$d(\hat{\mathbf{x}}_{i \rightarrow j}(\alpha'), \mathbf{x}_j) \leq d(\hat{\mathbf{x}}_{i \rightarrow j}(\alpha), \mathbf{x}_j)$$

3. **Smoothness:** The interpolation function $\hat{\mathbf{x}}_{i \rightarrow j}(\alpha)$ is Lipschitz continuous with a constant K :

$$\|\hat{\mathbf{x}}_{i \rightarrow j}(\alpha), \hat{\mathbf{x}}_{i \rightarrow j}(\alpha + t)\| \leq K|t|$$

4. **Credibility:** $\forall \alpha \in [0, 1]$ We require that it is highly probable that interpolated images, $\hat{\mathbf{x}}_{i \rightarrow j}(\alpha)$ belong to \mathcal{X} . Namely,

$$P(\hat{\mathbf{x}}_{i \rightarrow j}(\alpha)) \geq 1 - \beta, \quad \text{for some constant } \beta \geq 0$$

2.1. Proposed Approach

Following the above definitions for an admissible interpolation, we propose a new approach, called **Autoencoder Adversarial Interpolation** (AEAI), which shapes the latent space according to the above requirements. The general architecture comprises a standard autoencoder with an encoder, $\mathbf{z} = f(\mathbf{x})$, and a decoder $\hat{\mathbf{x}} = g(\mathbf{z})$. We also train a discriminator $D(\mathbf{x})$ to differentiate between real and interpolated data points. For pairs of input data points $\mathbf{x}_i, \mathbf{x}_j$, we linearly interpolate between them in the latent space: $\mathbf{z}_{i \rightarrow j}(\alpha) = (1 - \alpha)\mathbf{z}_i + \alpha\mathbf{z}_j$, where $\alpha \in [0, 1]$. The first requirement is that we would like $\hat{\mathbf{x}}_{i \rightarrow j}(\alpha) = g(\mathbf{z}_{i \rightarrow j}(\alpha))$ to look real and fool the discriminator D . Additionally, we add a cycle-consistency loss that encourages the latent representation of $\hat{\mathbf{x}}_{i \rightarrow j}(\alpha)$ to be mapped back into $\mathbf{z}_{i \rightarrow j}(\alpha)$ again; namely, $\hat{\mathbf{z}}_{i \rightarrow j}(\alpha) = f(g(\mathbf{z}_{i \rightarrow j}(\alpha)))$ should be similar to $\mathbf{z}_{i \rightarrow j}(\alpha)$. Finally, we add a smoothness loss that

drives the linear parameterization to form a smooth interpolation. Putting everything together we define the loss $\mathcal{L}_{i \rightarrow j}$ between pairs \mathbf{x}_i and \mathbf{x}_j as follows:

$$\mathcal{L}^{i \rightarrow j} = \mathcal{L}_R^{i \rightarrow j} + \lambda_A \mathcal{L}_A^{i \rightarrow j} + \lambda_C \mathcal{L}_C^{i \rightarrow j} + \lambda_S \mathcal{L}_S^{i \rightarrow j} \quad (1)$$

where $\mathcal{L}_R, \mathcal{L}_A, \mathcal{L}_C, \mathcal{L}_S$ are the reconstruction, adversarial, cycle, and smoothness losses, respectively. The first term \mathcal{L}_R is a standard reconstruction loss and is calculated for the two endpoints \mathbf{x}_i and \mathbf{x}_j :

$$\mathcal{L}_R^{i \rightarrow j} = \mathcal{L}(\mathbf{x}_i, \hat{\mathbf{x}}_i) + \mathcal{L}(\mathbf{x}_j, \hat{\mathbf{x}}_j)$$

where $\mathcal{L}(\cdot, \cdot)$ is some loss function between the two images (we used the L_2 distance or the perceptual loss (Johnson et al., 2016)) and $\hat{\mathbf{x}}_k = g(f(\mathbf{x}_k))$. \mathcal{L}_A is the adversarial loss that encourages the network to fool the discriminator so that interpolated images are indistinguishable from the data in the target domain \mathcal{X} :

$$\mathcal{L}_A^{i \rightarrow j} = \sum_{n=0}^M -\log D(\hat{\mathbf{x}}_{i \rightarrow j}(n/M))$$

where $D(\mathbf{x}) \in [0, 1]$ is a discriminator trying to distinguish between images in the training set and the interpolated images. The cycle-consistency loss \mathcal{L}_C encourages the encoder and the decoder to produce a bijective mapping:

$$\mathcal{L}_C^{i \rightarrow j} = \sum_{n=0}^M \|\mathbf{z}_{i \rightarrow j}(n/M) - \hat{\mathbf{z}}_{i \rightarrow j}(n/M)\|^2$$

where $\hat{\mathbf{z}}_{i \rightarrow j}(\alpha) = f(g(\mathbf{z}_{i \rightarrow j}(\alpha)))$. The last term \mathcal{L}_S is the smoothness loss encouraging $\hat{\mathbf{x}}(\alpha)$ to produce smoothly varying interpolated points between \mathbf{x}_i and \mathbf{x}_j :

$$\mathcal{L}_S^{i \rightarrow j} = \sum_{n=0}^M \left\| \frac{\partial \hat{\mathbf{x}}_{i \rightarrow j}(\alpha)}{\partial \alpha} \right\|_{\alpha=n/M}^2$$

where $\|\partial \hat{\mathbf{x}}_{i \rightarrow j}(\alpha) / \partial \alpha\|_{\alpha=\alpha_0}^2$ means that the derivative is taken at $\alpha = \alpha_0$. The three losses $\mathcal{L}_A, \mathcal{L}_C$ and \mathcal{L}_S are accumulated over $M + 1$ sampled points, from $\alpha = 0/M$ up to $\alpha = M/M$. Finally, we sum the $\mathcal{L}^{i \rightarrow j}$ loss over many sampled pairs.

In the next section, we explain the motivation for each of the four losses comprising $\mathcal{L}^{i \rightarrow j}$ in Equation 1 and describe how these losses promote the four conditions defined in Section 2.

2.2. Justification for the proposed approach

Figure 3 illustrates the justification for introducing the four losses. As seen in Plot A in Figure 3, the images $\mathbf{x}_i, \mathbf{x}_j$, which lie on the data manifold in the image space (solid black curve), are mapped back to the original images thanks




Figure 3. Data interpolation using autoencoders. Two points x_i, x_j are located on the input data manifold (solid black line). The encoder $f(x)$ maps input points into the latent space z_i, z_j (red arrows). Linear interpolation in the latent space is represented by the blue dashed line. The interpolated latent codes are mapped back into the input space by the decoder $g(z)$ (blue arrows). See Section 2.2 for the contribution of each loss component for an admissible interpolation.

to the reconstruction loss $\mathcal{L}_R^{i \rightarrow j}$. This loss promotes the *boundary conditions* defined above. The reconstruction loss, however, is not enough as it neither directly affects in-between points in the image space nor the interpolated points in the latent space. Introducing the adversarial loss $\mathcal{L}_A^{i \rightarrow j}$ prompts the decoder $g(z_{i \rightarrow j}(\alpha))$ to map interpolated latent vectors back into the image manifold (Plot B). Considering the output of the discriminator $D(x)$ as the probability of image x to be in the target domain \mathcal{X} (namely, to be on the image manifold), the adversarial loss promotes the *credibility condition* defined above. As indicated in Plot B, the encoder $f(x)$ (red arrows) might, nevertheless, still map in-between images to latent vectors that are distant from the linear line in the latent space. Adding the cycle-consistency loss $\mathcal{L}_C^{i \rightarrow j}$ forces the reconstruction of interpolated latent vectors to be mapped back into the original vectors in the latent space (Plot C). The adversarial and cycle-consistency losses encourage bijective mapping (one-to-one and onto) between the input and the latent manifolds, while providing a realistic reconstruction of interpolated latent vectors. Lastly, the parameterization of the interpolated points, namely, $\alpha \in [0, 1]$, does not necessarily provide smooth interpolation in the image space (Plot C); constant velocity interpolation in the parameter α may not generate smooth transitions in the image space. The smoothness loss $\mathcal{L}_S^{i \rightarrow j}$ resolves this issue as it requires the distance between x_i and x_j to be evenly distributed along $\alpha \in [0, 1]$ (due to the L_2 norm). This loss fulfills the *smoothness condition* defined above (Plot D). If we consider the latent representation as a normed space representing the manifold distance $d(x_i, x_j) = \|z_i - z_j\|$, the linear interpolation in the latent space also satisfies the *monotonicity condition* defined above.

2.3. Implementation

The proposed architecture is visualized in Figure 4. At each iteration, we sample two images from our dataset. The two images (x_i, x_j) are encoded by the encoder f into (z_i, z_j) , respectively. We sample α uniformly be-

tween $[0, 1]$ and pass (α, z_i, z_j) to h , a non-learned layer, which calculates the linear interpolation in the latent space, namely, $z_{i \rightarrow j}(\alpha) = (1 - \alpha)z_i + \alpha z_j$. We then decode $z_{i \rightarrow j}(\alpha)$ and calculate the reconstruction loss $\mathcal{L}_R^{i \rightarrow j}$. Subsequently, we decode $z_{i \rightarrow j}(\alpha)$ and alternately provide the discriminator D with samples either from the training set or from $\hat{x}_{i \rightarrow j}(\alpha) = g(z_{i \rightarrow j}(\alpha))$. The discriminator is optimized using the standard GAN loss and is updated after every iteration. We then calculate the smoothness loss $\mathcal{L}_S^{i \rightarrow j}$ by taking the derivative of $\hat{x}_{i \rightarrow j}(\alpha)$ with respect to α . Finally, we pass $\hat{x}_{i \rightarrow j}(\alpha)$ through the encoder f to obtain $\hat{z}_{i \rightarrow j}(\alpha) = f(\hat{x}_{i \rightarrow j}(\alpha))$ for the cycle-consistency loss and add the loss $\mathcal{L}_C^{i \rightarrow j}(z_{i \rightarrow j}(\alpha), \hat{z}_{i \rightarrow j}(\alpha))$.

The chosen encoder architecture was VGG-inspired (Simonyan & Zisserman, 2014). We extract the features using convolutional blocks starting from 16 feature maps, gradually increasing the number of feature maps to reach 128 by the last convolutional block. We then flatten the extracted features and pass them through fully connected layers until we reach our desired latent dimensionality. The decoder architecture is symmetrical to that of the encoder. We use max-pooling after each convolutional block and batch normalization with ReLU activations after each learned layer. A random 80%-20% training-testing split was chosen for all experiments, using the same batch size and total number of examples in the dataset. During log grid-search hyperparameter optimization, we found that $\lambda_A = \lambda_C = 10^{-2}$ and $\lambda_S = 10^{-1}$ produce the best results. All experiments were performed using a single NVIDIA V100 GPU.

3. Related Work

In its simplest version, the autoencoder (Goodfellow et al., 2016; Kramer, 1991) is trained to obtain a reduced representation of the input, removing data redundancies while revealing the underlying factors of the data set. The reduced space, namely, the latent space, can be viewed as a ‘useful’ representation space in which data interpolation can be attempted. Many autoencoder improvements have

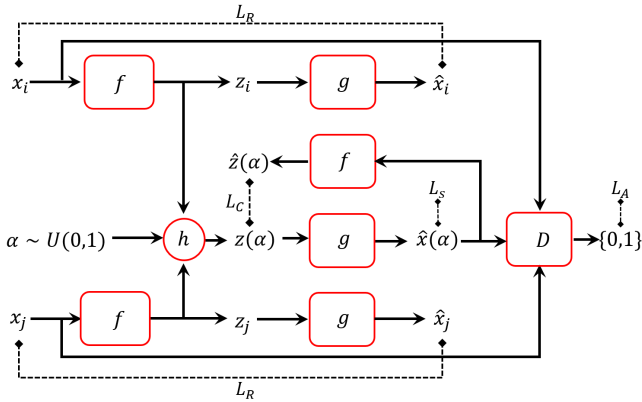


Figure 4. Our proposed architecture. Dotted lines represent the loss functions. h is a non-learned layer that performs latent linear interpolation.

been proposed in recent years, including new techniques designed for improved convergence and accuracy. Among these are the introduction of new regularization terms, new loss objectives (such as adversarial loss) and new network designs (Doersch, 2016; Kingma & Welling, 2013; Larsen et al., 2015; Makhzani et al., 2015; Vincent et al., 2010; Goodfellow et al., 2016; Vincent et al., 2010). Other new autoencoder techniques provide frameworks that attempt to shape the latent space to be efficient with respect to factor disentanglement or to make it conducive to latent space interpolation (Kingma & Welling, 2013; Bouchacourt et al., 2017; Vincent et al., 2008; Yeh et al., 2016; Higgins et al., 2016).

Within this second category, the variational autoencoder (VAE) and its derivatives were shown to be very successful in applying interpolation in the latent space, in particular for multimodal distributions, such as MNIST. The KL term in the VAE loss tends to cluster the modes in the latent space close to each other (Dieng et al., 2019). Consequently, linearly interpolating between different modes in the latent space may provide pleasing results that smoothly transition between the modes. Unfortunately, this cannot be applied to data points whose generating factors are continuous (in contrast to multimodal distributions) given that the KL loss term tends to fold the manifold tightly into a compact space making it highly non-convex.

(Berthelot et al., 2018) propose using a critic network to predict the interpolation parameter $\alpha \in [0, 1]$ while an autoencoder is trained to fool the critic. The motivation behind this approach is that the interpolation parameter α can be estimated for badly-interpolated images, while it is unpredictable for faithful interpolation. While this approach might work for multimodal data, it does not seem to work for data sampled from a continuous manifold. In such cases, the artifacts and the unrealistic-generated data do not provide any hint about the interpolating factor.

The GAIA method of (Sainburg et al., 2018) uses BEGAN

ion by Shaping the Latent Space

architecture composed of a generator and a discriminator, both based on autoencoders. The discriminator is trained to minimize the pixel-wise loss of real data and to maximize the pixel-wise loss of generated data (including interpolations). On the other hand, the generator is trained to minimize the loss of the discriminator for the interpolated data. The GAIA method is devoted to synthesizing realistic-looking images while ignoring the objective of image diversity and the need for smooth transitions between data points.

Perhaps the method most similar to our approach is the *adversarial mixup resynthesis* (AMR) of (Beckham et al., 2019). With the AMR method, a decoded mixup of latent codes $Mix(z_i, z_j)$ are encouraged to be indistinguishable from real samples by fooling a trained discriminator. This is similar to the adversarial loss introduced in our framework. Nevertheless, as elaborated in Section 2.2 and illustrated in Figure 3 (Plot B), the adversarial loss alone only amounts to generating realistic-looking interpolations, where the latent space is prone to mode collapse and sharp transitions along the interpolation paths.

In contrast to these methods, our additional smoothness and cycle-consistency requirements not only generate smooth transitions between data points but also ensure a diverse generation of realistic-looking images while avoiding mode collapse and abrupt transitions along the interpolating paths. This characteristic will be demonstrated in Sections 4 and 5.

Another approach is to separately learn the data manifold followed by performing interpolation on the manifold itself. This interpolation approach enforces points on the path to remain on the manifold, while using the least amount of change necessary (Shao et al., 2018). Common manifold learning techniques suggest regularizing the metric tensor (Chen et al., 2020) or preserving the local geometry of data points (Mishne et al., 2019). Although the latent interpolation scheme can assist autoencoders based manifold learning, those techniques do not take realistic interpolation as an explicit goal.

4. Results

Evaluating the realism of interpolation is often elusive. In the unsupervised scenario, where the ground-truth parameterization is unavailable, defining a path between two points p_i, p_j in the parameter space depends on the parameterization of the underlying factors governing the data, which is unknown. For example, in our synthetic pole dataset, the parameter space is (θ, ϕ) and there are infinitely many possible paths between any two points in that space, each of which can yield an admissible interpolation. Nevertheless, we evaluate the interpolation faithfulness both

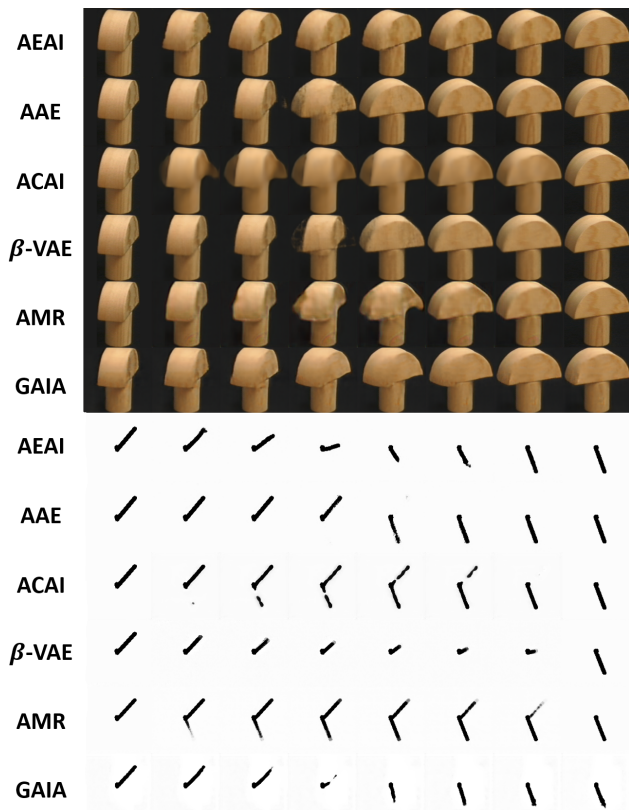


Figure 5. Each of the six rows presents linear interpolation of images from COIL-100 and our synthetic dataset for each of the methods tested.

qualitatively and quantitatively on various datasets based on the conditions we defined in Section 2.

4.1. Dataset

We tested our method against two different datasets: the synthetic pole dataset, which was rendered using the *Unity* game engine, where all images were taken from a fixed viewing position (from above) and the COIL-100 dataset. For the first dataset, a single illumination source was rotated at intervals of 5 degrees along the azimuth at different altitudes, ranging from 45 to 80 degrees with respect to the plane in 5-degree intervals. This dataset contains a total of 576 images. In the second dataset, to test our method against real images with complex geometric and photometric parameterization, we used the COIL-100 dataset (Nene et al., 1996) containing color images of 100 objects. The objects were placed on a motorized turntable against a black background. The images were taken at intervals of 5 degrees resulting in a total of 72 images for each class. All architectures were trained using a dataset that contained a single class from COIL-100 or the synthetic pole dataset.

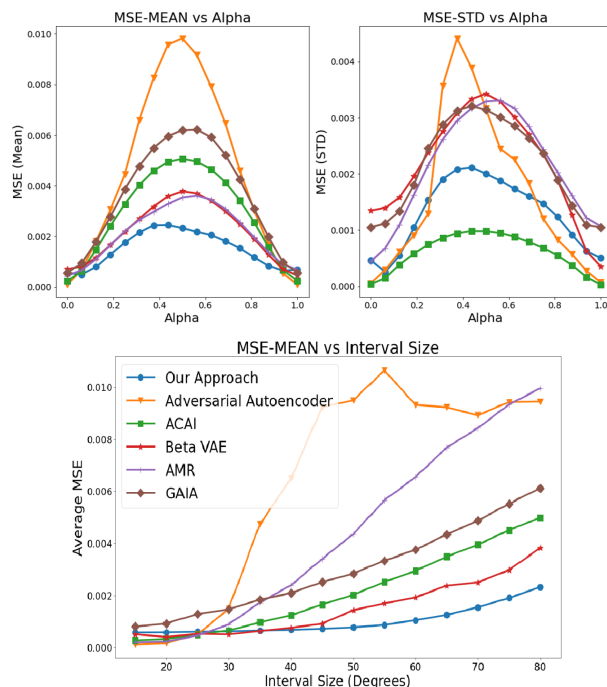


Figure 6. We use the parameterization of the dataset to evaluate the reconstruction accuracy of the AAE, ACAI, β -VAE, AMR, GAIA and our proposed method. Upper left graph: Averaged MSE vs. α values. Upper right graph: STD of MSE vs. α values. Lower graph: Averaged MSE of the interpolated images vs. the interval length.

4.2. Qualitative Assessments

Each one of the six rows in Figure 5 presents a linear interpolation of a single object from the COIL-100 dataset (upper part) and our pole dataset (lower part) sampled from the validation dataset. We compared the results of the β -Variational Autoencoder (β -VAE) (Higgins et al., 2016), the Adversarial Autoencoder (AAE) (Makhzani et al., 2015), the Adversarially Constrained Autoencoder Interpolation (ACAI) (Berthelot et al., 2018), the Generative Adversarial Interpolative Autoencoding method (GAIA) (Sainburg et al., 2018), the Adversarial Mixup Resynthesis (AMR) (Beckham et al., 2019) and our approach—Autoencoder Adversarial Interpolation (AEAI). In the experiments with both datasets, we used a latent dimensionality of 256. From Figure 5 it can be seen that our proposed method provides realistic-looking reconstructions and an admissible interpolation between modes. The AAE and β -VAE interpolations change abruptly between modes and introduce small artifacts during reconstruction. The ACAI produces unrealistic transitions and artifacts during reconstruction, especially in the mid-range of the α -values. The GAIA method produces realistic transitions with small artifacts during reconstruction while the AMR produces significant artifacts and gradual transitions between modes. Additional qualitative

results are presented in the supplementary material.

4.3. Quantitative Assessments

For a quantitative comparison we used the COIL-100 dataset. We fixed an interval length, which is a multiplicative of 5 degrees, and calculated the reconstruction error (MSE) against the available ground-truth images. We used an interval length of 80 degrees that resulted in 14 in-between images of a single object. The reconstruction error of the interpolated images is presented in Figure 6. Clearly, our method reduces the mean MSE and the standard deviation of the MSE for different alpha values. We then inspected the average reconstruction error on multiple intervals ranging from 15 to 80 degrees as presented in the right part of Figure 6. Note that our proposed method is able to reduce the reconstruction error of interpolated images consistently even when the interval length increases.




Figure 7. Predicting the interpolated alpha value based on the L_2 distance of the interpolated image to the closest image in the dataset. The dots represent the median and the colored area corresponds to the interquartile range.

To assess the transition smoothness from one sample to the other, we compared each interpolated image $\hat{x}_{i \rightarrow j}(\alpha)$ to the closest image in the dataset in terms of the L_2 distance and assigned the alpha value for the interpolated image according to the retrieved image. We repeated this process for all the intervals of length 70. Figure 7 presents the scatter diagrams for each method. It is demonstrated that our framework consistently retrieves the best value of alpha with a smaller interquartile range (IQR). The next experiment was applied to the synthesized pole dataset. As above, we retrieved the closest image in terms of MSE in the image space, and measured the L_2 distance in the *parameter*




Figure 8. We sampled two images x_i, x_j and linearly interpolated between them in latent space. For each interpolated image, we retrieved the closest image in terms of MSE from the dataset. The blue and orange lines present the averaged L_2 distance, in the parameter space (θ, ϕ) , between the retrieved image and x_i, x_j , respectively. The red lines represent perfect interpolation smoothness.

space between the interpolated image and the source image ($\alpha = 0$) and between the interpolated image and the target image ($\alpha = 1$). We repeated this process on multiple intervals of different lengths on both θ and ϕ , and present the average distance from the source and target images as a function of the interpolation variable, α . Figure 8 shows the results for each tested method. It is demonstrated that all methods exhibit monotonicity characteristics; however our approach outperforms the other methods with respect to the smoothness of the parameterization. Our approach tightly follows the linear parameterization due to the explicit smoothness term we incorporated into the training scheme.

5. Ablation Study

We present an ablation study of our unsupervised interpolation framework presented in Figure 4 above. As seen in the lower left part of Figure 9, without a significant contribution from the discriminator, the reconstructed images resulting from interpolating latent vectors are unrealistic and exhibit severe artifacts and non-smooth transition between modes. Without the cycle-consistency loss, interpolated images are relatively realistic; however, they change modes abruptly, exhibiting artifacts when transitioning from one mode to another. Adding both cycle consistency and the discriminator results in realistic transitions from mode to mode as can be seen in the bottom right part of Figure 9. Nevertheless, there are instances of consecutive interpolated images that show little to no change as can be seen in the two leftmost columns of the bottom right part of Figure 9. When introducing the smoothness loss, as can be seen in the upper left

part of Figure 9, we get both smooth transitions and realistic reconstructions of interpolated latent vectors.




Figure 9. Each of the four blocks presents bilinear interpolation of four ground truth images that reside in each corner of the block. Upper left: Bilinear interpolation results of our approach with all loss components. Upper right: Removing the cycle-consistency contribution from the loss function. Lower left: Removing the discriminator contribution. Lower right: Removing the smoothing contribution.

We present a quantitative analysis of our ablation study in Figure 10. For each case, we demonstrate the average MSE error between the interpolated image and ground truth images retrieved from our dataset. In the upper part of Figure 10, we fix an interval length of 80 degrees and iterate over all such intervals in our dataset. We split the interval into 16 images separated by 5 degrees and obtain 14 in-between images. For every interpolated image we retrieve the corresponding ground truth image in our dataset and present the average MSE and standard deviation on all such intervals. In the lower part of Figure 10 we repeat this process on multiple interval sizes ranging from 15 degrees to 80 degrees. We show that each element in our loss function 1 contributes to reducing the mean and variance of the reconstruction error. We further support our justification presented in Figure 3 by showing that the largest contribution to the reduction of the MSE stems from the introduction of the discriminator, which keeps the reconstruction of interpolated images consistent with the dataset. The addition of the cycle-consistency and smoothness losses further improves the results by encouraging a smooth bijective mapping. Additional results are available in the supplementary material.

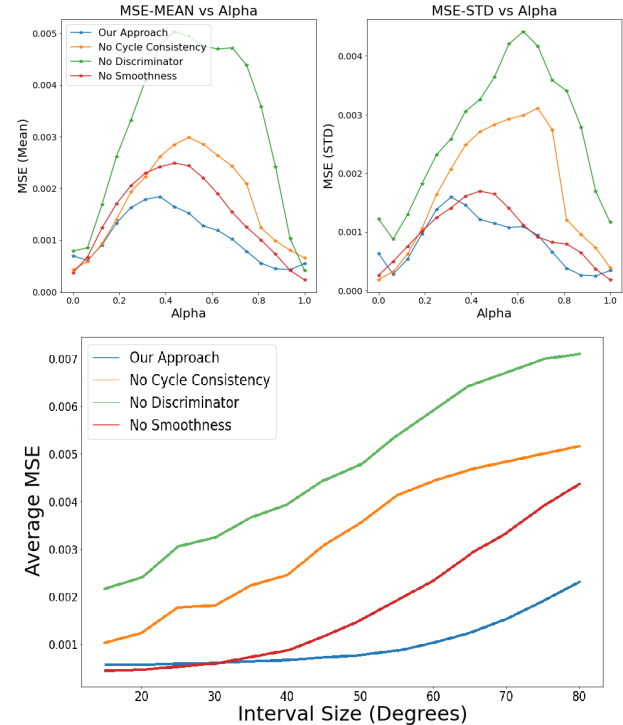


Figure 10. Ablation study of the reconstruction accuracy. Upper left graph: Average MSE vs. α values. Upper right graph: STD of MSE vs. α values. Lower graph: Averaged MSE of the interpolated images vs. the interval length.

6. Conclusion & Discussion

The problem of realistic and faithful interpolation in the latent spaces of generative models has been tackled successfully in the last few years. Nevertheless, it is our opinion that generative approaches that deal with manifold data are not as common as multimodal data, and this misinterpretation of manifold data harms the competence of generative models to deal with them successfully. In this work, we argue that the manifold structures of data generated from continuous factors should be taken into account. Our main contribution is applying convexity regularization using adversarial and cycle-consistency losses. Applying this technique on small datasets of images, taken from various viewing conditions, we managed to greatly improve the fidelity of interpolated images. We also implemented a smoothness loss and improved the non-uniform parameterization of the latent manifold. In future work, we intend to further investigate properties of latent manifolds, in particular, capable of generating admissible interpolation between both categorized and continuous data, and use the proposed approach as a general regularizer method for generative models with few training examples.

Acknowledgements

This work was supported by the Israeli Ministry of Science and Technology under The National Foundation for Applied Science (MIA).

References

- Beckham, C., Honari, S., Verma, V., Lamb, A. M., Ghadiri, F., Hjelm, R. D., Bengio, Y., and Pal, C. On adversarial mixup resynthesis. In *Advances in neural information processing systems*, pp. 4346–4357, 2019.
- Bengio, Y., Courville, A., and Vincent, P. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8), 2013.
- Berthelot, D., Raffel, C., Roy, A., and Goodfellow, I. Understanding and improving interpolation in autoencoders via an adversarial regularizer. *arXiv preprint arXiv:1807.07543*, 2018.
- Bouchacourt, D., Tomioka, R., and Nowozin, S. Multi-level variational autoencoder: Learning disentangled representations from grouped observations. *arXiv preprint arXiv:1705.08841*, 2017.
- Chen, N., Klushyn, A., Ferroni, F., Bayer, J., and Van Der Smagt, P. Learning flat latent manifolds with VAEs. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1587–1596. PMLR, 13–18 Jul 2020.
- Dieng, A. B., Kim, Y., Rush, A. M., and Blei, D. M. Avoiding latent variable collapse with generative skip models. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2397–2405. PMLR, 2019.
- Doersch, C. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.
- Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. *Deep learning*. MIT press Cambridge, 2016.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vaes: Learning basic visual concepts with a constrained variational framework. *Arxiv*, 2016.
- Johnson, J., Alahi, A., and Fei-Fei, L. Perceptual losses for real-time style transfer and super-resolution. In *The European Conference on Computer Vision (ECCV)*, 2016.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Kramer, M. A. Nonlinear principal component analysis using autoassociative neural networks. *AIChE journal*, 37(2):233–243, 1991.
- Larsen, A. B. L., Sønderby, S. K., Larochelle, H., and Winther, O. Autoencoding beyond pixels using a learned similarity metric. *arXiv preprint arXiv:1512.09300*, 2015.
- Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., and Frey, B. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- Mathieu, M. F., Zhao, J. J., Zhao, J., Ramesh, A., Sprechmann, P., and LeCun, Y. Disentangling factors of variation in deep representation using adversarial training. In *Advances in neural information processing systems*, pp. 5040–5048, 2016.
- Mishne, G., Shaham, U., Cloninger, A., and Cohen, I. Diffusion nets. *Applied and Computational Harmonic Analysis*, 47(2):259–285, 2019. ISSN 1063-5203. doi: <https://doi.org/10.1016/j.acha.2017.08.007>. URL <https://www.sciencedirect.com/science/article/pii/S1063520317300957>.
- Nene, S. A., Nayar, S. K., and Murase, H. object image library (coil-100). Technical report, 1996.
- Sainburg, T., Thielk, M., Theilman, B., Migliori, B., and Gentner, T. Generative adversarial interpolative autoencoding: adversarial training on latent space interpolations encourage convex latent distributions. *arXiv preprint arXiv:1807.06650*, 2018.
- Shao, H., Kumar, A., and Fletcher, P. T. The riemannian geometry of deep generative models. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 428–4288, 2018. doi: 10.1109/CVPRW.2018.00071.
- Shu, Z., Sahasrabudhe, M., Alp Guler, R., Samaras, D., Paragios, N., and Kokkinos, I. Deforming autoencoders: Unsupervised disentangling of shape and appearance. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Verma, V., Lamb, A., Beckham, C., Courville, A., Mitliagkis, I., and Bengio, Y. Manifold mixup: Encouraging meaningful on-manifold interpolation as a regularizer. *arXiv preprint arXiv:1806.05236*, 2018.
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international*

conference on Machine learning, pp. 1096–1103. ACM, 2008.

Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P.-A. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(Dec):3371–3408, 2010.

Yeh, R., Liu, Z., Goldman, D. B., and Agarwala, A. Semantic facial expression editing using autoencoded flow. *arXiv preprint arXiv:1611.09961*, 2016.

Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018.