

Towards Photographic Image Manipulation with Balanced Growing of Generative Autoencoders

Ari Heljakka^{1,2} Arno Solin¹ Juho Kannala¹

¹Department of Computer Science, Aalto University, Espoo, Finland

²GenMind Ltd., Espoo, Finland

firstname.lastname@aalto.fi

Abstract

We present a generative autoencoder that provides fast encoding, faithful reconstructions (e.g. retaining the identity of a face), sharp generated/reconstructed samples in high resolutions, and a well-structured latent space that supports semantic manipulation of the inputs. There are no current autoencoder or GAN models that satisfactorily achieve all of these. We build on the progressively growing autoencoder model PIONEER, for which we completely alter the training dynamics based on a careful analysis of recently introduced normalization schemes. We show significantly improved visual and quantitative results for face identity conservation in CELEBA-HQ. Our model achieves state-of-the-art disentanglement of latent space, both quantitatively and via realistic image attribute manipulations. On the LSUN Bedrooms dataset, we improve the disentanglement performance of the vanilla PIONEER, despite having a simpler model. Overall, our results indicate that the PIONEER networks provide a way towards photorealistic face manipulation.

1. Introduction

Advances in generative image modelling with deep neural networks have raised expectations for delivering new tools for photographic image manipulation and exploration. Generative image modelling involves training the model by a wide variety of data, such as pictures of faces or bedrooms, and allowing it to learn features and structure present in the data. Recent works (e.g. [19, 20, 3]) on using Generative Adversarial Networks (GANs, [10]) for images of high quality and resolution have showed that *generating* random beautiful and sharp high-resolution images is achievable by current tools—given a fair deal of engineering and compute.

However, to manipulate existing images and other content, we also need inference, not only the capability to generate random samples. A model that can encode the sample of interest into a well-structured latent space allows us to manip-




Figure 1. Six-way example interpolations in the latent space between reconstructions of previously unseen input images (bottom) at 256×256 resolution. Best viewed zoomed in.

ulate the sample via that latent representation. In absence of extensions, a GAN has no inference component. In contrast, generative autoencoder models allow both generation and reconstruction, with a bi-directional mapping between the latent feature space and image space. The latent space can be exploited by interpolating between two or more points that represent images (see Fig. 1), by finding directions that encode an individual feature of interest (see Fig. 6) to modify new images, and for other downstream tasks.

The state-of-the-art in this direction was set by three papers simultaneously published in late 2018: the flow-based GLOW method [25], the introspective variational autoen-



Figure 2. Example of reconstruction quality in 256×256 resolution with typical images from the CELEBA-HQ test set (top row), by our balanced PIONEER (middle) and baseline PIONEER (bottom). Here, the input images are encoded into 512-dimensional latent feature vector and decoded back to the original dimensionality (middle and bottom rows). The encoding–decoding of balanced PIONEER tends to preserve facial features, orientation, expressions, and hair style. Small mistakes can still be observed, especially in male subjects. (Features highly underrepresented in the limited CELEBA-HQ training set, such as a dark-skinned face with atypical orientation, will result in more pronounced failures.)

coder (IntroVAE, [16]), and the progressively growing generative autoencoder (PIONEER, [13]) based on the Adversarial Generator–Encoder (AGE, [40]). GLOW offers tractable likelihood estimates and a degree of attribute manipulation but ultimately suffers from mode collapse issues. IntroVAE offers high sample quality but with *e.g.* human faces, it often only retains overall *topology*, not the identity [16].

Here, we build on the progressively growing generative autoencoder concept, primarily the PIONEER. We seek to balance the training in higher resolutions—a common challenge for generative models (see *e.g.* discussion in [16]). PIONEER was previously shown to reconstruct images only up to 128×128 . Our model provides 256×256 reconstructions (Fig. 2), generates realistic and diverse samples, and learns disentangled latent representations of features (Fig. 6).

The contributions and results of this paper are as follows. (i) We show that an AGE-based autoencoder model (balanced PIONEER) can learn a high-resolution image dataset, combining all the strengths of both autoencoder and GAN models—fast encoding, faithful reconstruction of inputs, and sharp sample generation (contrary to what is implied in [16]). We also show state-of-the-art disentanglement capabilities in terms of Perceptual Path Length metric [20], along with image attribute modifications at higher resolutions than pre-

viously shown for unsupervised general-purpose models.

(ii) We propose a modified PIONEER model with completely altered training dynamics, considerably improving the learning capacity. When reconstructing face images of CELEBA-HQ, the baseline PIONEER often loses the identity. The cause is related to a large dynamic range and occasional collapse of the encoder–decoder minimax game, complicated by three different weight normalization schemes: Pixel Norm (PN), Equalized Learning Rate (EQLR) and Spectral Normalization. For the first time, we show systematic comparison of the effects of each on a generative model. By improving the loss function with a margin term, we are able to drop PN and EQLR while solving the stability problem.

(iii) With these improvements, our model can reconstruct 256×256 CELEBA-HQ images with substantially better conservation of identity and attribute editability than shown for the baseline models such as the state-of-the-art model IntroVAE, with almost comparable sample quality. In comparison to the vanilla PIONEER, we show substantially better conservation of face identity (via L2 distance in the embedding space of a face recognition model [22, 9]) and significantly improved LPIPS [44] and FID [15] scores. In LSUN Bedrooms dataset, we improve the disentanglement performance of the vanilla PIONEER despite our model being simpler.

2. Related work

Our work builds upon the previous research in generative image models, such as variational autoencoders (VAEs, [23, 18]), autoregressive models, flow models, and GAN variants. Our approach borrows ideas from both autoencoders and GANs. In this section, we give a brief overview of the background literature and conclude with discussion about the closest related works.

The basic idea of a GAN is to train two networks, generator and discriminator, in a competitive manner so that (a) the generator learns to produce images from the same distribution as the training data and (b) the discriminator learns to distinguish the synthetic images produced by the generator from the real training samples as well as possible. If successful, the generator has become good enough so that the discriminator can no longer make the distinction. The generator starts from a compact random latent code.

During the recent years there has been rapid progress related to GAN-based models and applications. Many recent improvements improve the stability and robustness of the training process by suggesting new loss functions [1], regularization methods [33, 35, 31], multi-resolution training [19, 43], architectures [20], or combinations of these.

However, despite the notable progress in image generation, it is widely accepted that the capability for realistic image synthesis alone is not sufficient for most applications, such as image manipulation, where we start from an existing image. This calls for a model with an encoder, such as a VAE. On the other hand, autoencoders tend to suffer from blurry outputs, and are often used for only low resolution images (even in recent works such as *e.g.* [37, 38, 5], see also comparison in [13]).

Thus, there have been many efforts to combine GANs with autoencoder models, (*e.g.* [4, 34, 7, 8]). For instance, [7] and [8] proposed utilizing three deep networks in order to learn functions that enable mappings between the data space and the latent space in both directions. That is, besides the typical autoencoder architecture, consisting of a decoder (*i.e.* generator) and encoder networks, their approach uses an additional discriminator network, which is trained to classify tuples of image samples with their latent codes. Other authors introduce additional discriminator networks besides the generator and encoder. For example, [28, 4] use a GAN-like discriminator in sample space and [30, 32] in latent space. Nevertheless, the image synthesis performance of these hybrid models has not yet been shown to match state-of-the-art of purely generative models [19, 20].

In this paper, we build upon PIONEER [13], based on the adversarial generator–encoder (AGE) [40]. In contrast to many other previous works, these two models consist of only two deep networks, a generator and an encoder, which represent the mappings between the image space and latent space. In addition, the method of progressive network

growing, adapted from [19], is utilized in [13].

The results of [13] are promising, and both synthesis and reconstruction have good quality in relatively high image resolutions. However, in this paper, we show that [13] suffers from large fluctuations of the competing divergence terms of the adversarial loss, and this seems to hamper optimization and convergence thereby limiting performance.

Besides [13], other recent and related works are IntroVAE [16] and GLOW [25]. Based on VAE, IntroVAE is fundamentally different from PIONEER that is based on AGE. IntroVAE has been shown to produce high quality samples, but not to faithfully conserve sample details such as identity of faces (see Fig. 3 in [16]). Mere conservation of overall image topology is generally insufficient for manipulating image attributes. Based on the contributions of our paper, we are able to show that AGE-based generative models are capable of producing competitive results at 256×256 resolution, while conserving the face identity better than IntroVAE. This is in contrast to the observations in [16], where the authors were not able to make AGE training converge with large image resolutions. This finding is particularly promising since the PIONEER model has a simpler yet more powerful architecture than the corresponding purely generative model PGGAN [19]. It shows that the conventional GAN paradigm of a separate discriminator network is not necessary for learning to infer and generate image data sets.

Our model learns to manipulate image attributes in a fully unsupervised manner, in contrast to supervised approaches where the class information is provided during training (*e.g.* [27, 14]), and to models only capable of specific discrete domain transformations [46, 21, 41, 6, 17]. In practice, all prior unsupervised work uses 64×64 resolution, such as every model cited in a recent large-scale autoencoder comparison [29]. In the GAN research line, the state-of-the-art models such as [20] could be used for this in high resolution, but they have no encoder to deal with new input images.

3. Methods

In this section, we start from the basics of PIONEER [13] which in turn builds on [40] and [19] (Sec. 3.1). From the perspective of image manipulation in 256×256 resolution (and above), the regular PIONEER reaches reasonable sample quality and diversity, but the reconstructions are often not faithful to the originals (unlike in 128×128 CELEBA).

Therefore, our main task in this paper is to substantially improve the reconstructions without sacrificing other performance aspects. For this, we must address two issues: (i) On the training behavior level, the PIONEER training dynamics are difficult to optimize, due to wildly oscillating Kullback–Leibler divergence (KL) terms during training. Further, occasionally the encoder and decoder completely diverge and the training performance collapses (Fig. 3a). (ii) On the algorithmic level, the three different normalization schemes

used in baseline PIONEER compound the difficulty of understanding the behavior of the model, and together appear to contribute to the issue (1) in complicated ways. Instead of proposing yet another network architecture, we seek to find the minimum sufficient changes to simplify and stabilise the training workflow.

We benefit from first simplifying the algorithm, and reduce the three normalization schemes to only one (Sec. 3.2). This simplification comes at the cost of even less stable training dynamics. However, the behavior becomes more straightforward, and we are able to stabilize the training by bounding the difference between the competing KL terms, preventing them from diverging (Sec. 3.3). We also compare the impact of each combination of normalization strategies on early training stages (Fig. 3a).

Finally, we analyze the representational power of our improved model by measuring the degree of disentanglement in the model’s latent space and demonstrate realistic image feature manipulations (Sec. 3.4).

3.1. Model training dynamics

In PIONEER, the encoder ϕ and decoder θ are trained separately during each training step, with separate loss functions. Because of this separation and the presence of loss terms that are the exact opposites of each other, the training can be called adversarial. With $\mathbf{x} \sim X$ as the training samples and $\hat{\mathbf{x}} = \theta(\hat{\mathbf{z}})$ where $\hat{\mathbf{z}} \sim N(\mathbf{0}, \mathbf{I})$, the encoder loss consists of the encoder trying to push the distribution of the latent codes of the training samples $q_\phi(\mathbf{z} | \mathbf{x})$ towards a unit Gaussian distribution $N(\mathbf{0}, \mathbf{I})$ and the distribution of the codes of the generated samples $q_\phi(\mathbf{z} | \hat{\mathbf{x}})$ away from $N(\mathbf{0}, \mathbf{I})$. The decoder tries to do the opposite to the generated samples. Furthermore, the encoder attempts to minimize reconstruction error $L_{\mathcal{X}}$ with L1 distance in sample space \mathcal{X} , and the decoder to minimize reconstruction error $L_{\mathcal{Z}}$ with cosine distance in latent code space \mathcal{Z} . Hence, as explained in detail in [13], the full loss functions are as follows:

$$L_\phi = D_{\text{KL}}[q_\phi(\mathbf{z} | \mathbf{x}) \| N(\mathbf{0}, \mathbf{I})] - D_{\text{KL}}[q_\phi(\mathbf{z} | \hat{\mathbf{x}}) \| N(\mathbf{0}, \mathbf{I})] + \lambda_{\mathcal{X}} L_{\mathcal{X}}, \quad (1)$$

$$L_\theta = D_{\text{KL}}[q_\phi(\mathbf{z} | \hat{\mathbf{x}}) \| N(\mathbf{0}, \mathbf{I})] + \lambda_{\mathcal{Z}} L_{\mathcal{Z}}, \quad (2)$$

with D_{KL} denoting Kullback–Leibler divergence. The reconstruction loss terms are defined as:

$$L_{\mathcal{X}}(\theta, \phi) = \mathbb{E}_{\mathbf{x} \sim X} \|\mathbf{x} - \theta(\phi(\mathbf{x}))\|_1, \quad (3)$$

$$L_{\mathcal{Z}}(\theta, \phi) = \mathbb{E}_{\mathbf{z} \sim N(\mathbf{0}, \mathbf{I})} [1 - \mathbf{z}^\top \phi(\theta(\mathbf{z}))]. \quad (4)$$

Vectors in \mathcal{Z} are always normalized to unity. Contrasted to regular autoencoders, the second major aspect of PIONEER is the progressive growing of the network architecture during training, as follows. The encoder and decoder are divided into residual convolution/deconvolution blocks operating on

separate resolutions (16×16 , 32×32 , ...). The training is divided into phases during which we only train the blocks that operate up to that resolution. Once that resolution is sufficiently learnt, we gradually fade in the next level blocks, and so on. (For more details, see [13, 19].)

When training on resolutions 128×128 or higher, we observe that around the time when the image generation results begin to get worse, the encoder tends to assign increasingly low KL divergence estimates for the training samples, and increasingly large ones for the generated samples. In other words, the encoder wins the game. We will keep this observation in mind, and now turn to the weight normalization schemes that underlie the training dynamics.

3.2. Simplification of the normalization scheme

In [13], three unrelated implicit regularization techniques are used to stabilize the training: (i) equalized learning rate [19] and (ii) pixel norm [19] in the decoder, as well as (iii) spectral normalization [33] in the encoder. (i) is used to scale generator weights by dividing each weight with He’s initializer [12] at runtime. (ii) normalizes, for each convolutional layer, the feature vector of each pixel to unit length. In (iii), the spectral norm of each layer of the network is constrained directly during every computation pass by dividing each network weight matrix by its largest singular value. This keeps the Lipschitz constant under control. Intuitively, (iii) ensures that the weights do not amplify the ‘scale’ of the input signal, but allows them to freely ‘rotate’ the signal.

Since (i) scales the weights based on the number of input connections, it merely ensures that the overall scales of weight gradient updates are balanced against the changing layer sizes. While this helps to even out extreme cases, it does not guarantee Lipschitz continuity.

Method (ii) scales the outputs of the convolution operations. First, it prevents activations from escalating to large values. Second, it eliminates much of the information about how the activation of a given filter varies across locations. The filter might activate more strongly on one pixel than on another, but this difference may be ‘scaled away’ depending on the other filter activations. This explains the observed dampening effect that (i) and (ii) have on learning, with more stability at the cost of capacity.

Informally, we can construe the regularization as having two goals: (A) stable individual training steps and (B) preventing the encoder–decoder competition from spiralling out of control. (i) and (ii) meet (A) but fail (B). As their effect is local, they provide nothing to prevent the encoder from over-powering the decoder over time (Fig. 3). Likewise, (iii) guides only the local learning, but with less dampening.

As (iii) is more principled an approach than (i) and (ii), we could in fact use it to replace (i) and (ii) altogether. Hypothetically, without the strong dampening effect, this buys us more learning capacity. Empirically, however, it then leads

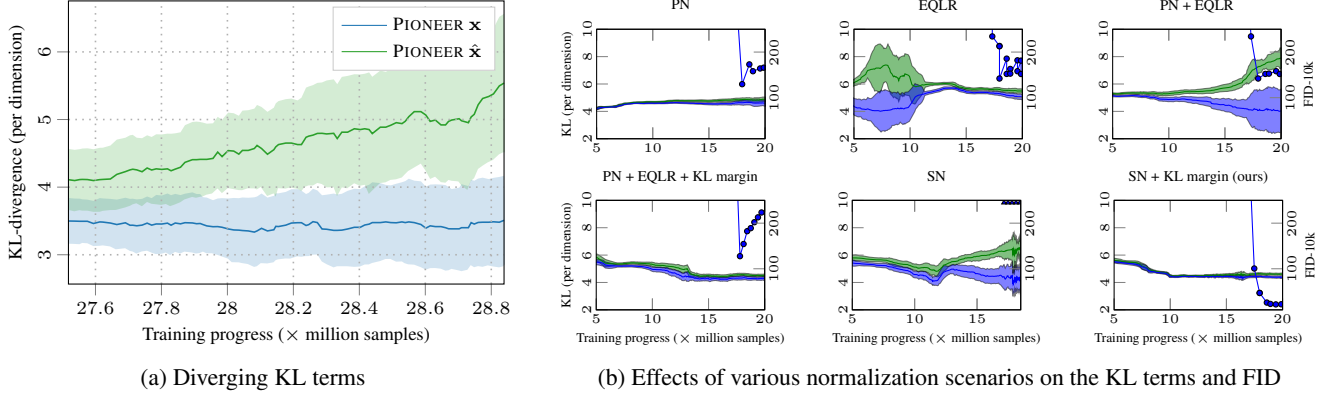


Figure 3. (a) The competing KL divergence terms that the encoder assigns for \mathbf{x} (training samples) and $\hat{\mathbf{x}}$ (generated samples) in PIONEER show two pathological properties. Relatively early on, the encoder overpowers the decoder leading the terms to diverge from each other, and the dynamic range of each grows, leading to an upper bound on the learning capacity. (b) The individual contribution of our changes, shown for KL divergence terms for early stages of CELEBA-HQ training. FID (blue dots) is shown for 128×128 stage onwards (with stage fade-in starting at 16.52M samples). PixelNorm (PN) or Equalized Learning Rate (EQLR) alone or in combination (top) will lead to poor FID even with KL margin (bottom left). On the other hand, merely adding Spectral Normalization (SN) without KL margin will lead to divergence of the terms and training collapse (bottom middle). Adding SN with the margin produces the optimal outcome (bottom right).

to unbalanced competition even faster than with (i) and (ii).

Now, the key insight here is to use the implicit regularizer (iii) to address goal (A) *only*. To this end, the weaker constraints of (iii) will suffice for encoder and decoder alike, allowing us to discard (i) and (ii) altogether. For (B), we will proceed to explicitly improve the loss function, instead.

3.3. Competing KL divergence terms

We now return to the KL terms. In general, their absolute values do not seem to always correlate with better generation results. Instead, the difference between them is critical.

We thus set out to regularize this part of the training by adding a simple hinge loss to limit the gradient reward for the encoder. A major component of the gradient of the encoder comes from the gap between the KL terms. We wish to ensure that the encoder is not motivated to increase the gap too much, so we define a single margin term M_{gap} that defines the upper bound of the gap, and then modify Eq. (1) as follows into a hinge-loss form:

$$L_{\phi} = \max(-M_{\text{gap}}, D_{\text{KL}}[q_{\phi}(\mathbf{z} | \mathbf{x}) \| \mathcal{N}(\mathbf{0}, \mathbf{I})] - D_{\text{KL}}[q_{\phi}(\mathbf{z} | \hat{\mathbf{x}}) \| \mathcal{N}(\mathbf{0}, \mathbf{I})]) + \lambda_{\mathcal{X}} L_{\mathcal{X}}. \quad (5)$$

Because the latter KL term is, in practice, almost always larger than the former, the gap is negative and needs to be bounded from below. Now, Eq. (5) alone would not care about the absolute values of the KL terms at all. However, since Eq. (2) remains unchanged, the decoder training provides the force that drives the $D_{\text{KL}}[q_{\phi}(\mathbf{z} | \hat{\mathbf{x}}) \| \mathcal{N}(\mathbf{0}, \mathbf{I})]$ lower. Combined with Eq. (5), the result is a force that pushes the other KL term lower, too (see Fig. 3 for an example of diverging KLs and the effect of bounding the gap).

After applying this change, the training becomes stable again, regardless of resolution, while using our simplified weight normalization scheme. There is no collapse of the training, and furthermore, we tend to see steady improvements well beyond the capacity of the original PIONEER. With the approach in Eq. (5), we do not need to try to reduce the learning rate to dampen the gradients. We note that the large gradients are not a problem in the early pre-training stages. On the contrary, applying a heavy-handed margin in the early stages will reduce the rate of learning too much. Hence, it is sufficient to only apply the margin after the pre-training. Alternatively, we could have defined the margin to be dependent on the progressive stage of the training, but more experiments should be done to confirm the generality of such margin choices.

3.4. Disentanglement of latent space

One of many ways to measure disentanglement of the latent space is to find latent directions that correspond to specific factors of variation in the sample space. In an unsupervised training setup, these factors are unknown *a priori*. However, as a proxy measure, we can measure how smoothly the generated samples change while we move around in the latent space. This requires a good metric (inductive bias) for measuring such changes between images, such as LPIPS [44]. Following [20], we compute the Perceptual Path Length (PPL) by repeatedly taking a short constant-length random vector in the latent space, generating images at its endpoints, and measuring LPIPS between the points.

Visually, we can modify isolated image attributes, such as the degree of smiling, by finding the corresponding latent code vectors after the training. We take our existing model *trained without labels*, a set A of images with the desired




Figure 4. FID-10k values of CELEBA-HQ during training, with regular PIONEER and Balanced PIONEER. The FID measures the quality and diversity of generated images (smaller is better). The regular PIONEER scores are markedly worse and more volatile.

attribute, and a set B without it. *Requiring no additional training or optimization*, we take the difference between the mean latent code of each set. The resulting 512-dimensional difference vector can then be added to the code of any new real test image \mathbf{x}_{noF} , scaled to the desired intensity λ . The decoded image then gains (or, with negative λ , loses) the attribute (see Fig. 6):

$$\mathbf{x}_F(A, B) = \theta(\phi(\mathbf{x}_{\text{noF}}) + \lambda[\mathbb{E}_{\mathbf{x} \sim A} \|\phi(\mathbf{x})\| - \mathbb{E}_{\mathbf{x} \sim B} \|\phi(\mathbf{x})\|]).$$

4. Experiments

Any generative models can create data samples, but encoder-decoder models can also take input samples and modify them via latent space. Correspondingly, one can test such a model by reconstructing new test samples and interpolating between them in the latent space, evaluating some generated random samples, and analyzing the latent space.

The various ways of measuring generative models constitute an active field of research as such. To evaluate random samples, we look for both quality and diversity of the sampled distribution, typically via Fréchet inception distance (FID, [15]) or Inception Distance [36]. They come with various shortcomings [2] but as long as we use identical sample size, FID remains a reasonably reliable tool for comparing different models. For easy comparison to prior work, we measure against the training set (but see [26] for limitations).

To measure the extent to which face identity of an input image is conserved in the reconstructed face images, we used L2 distance between the 128-D feature embeddings of the two images, in the embedding space of a pre-trained DLib face recognition model [22, 9]. To measure the conservation of overall image features, we use LPIPS [44] which correlates with human judgement more closely than, for example, the commonly used pixel space L2 distance or Structural Similarity measures [45]. Finally, to evaluate the quality of

Table 1. Comparison of Fréchet Inception Distance (FID) and perceptual path length (PPL) of 256×256 images on CELEBA-HQ and LSUN dataset between Balanced PIONEER, regular PIONEER, PGGAN, IntroVAE and GLOW. Balanced PIONEER has the best PPL result, PGGAN has the best FID results. Improvement of the Balanced PIONEER over the regular PIONEER is clear on CELEBA-HQ. [16] provides the 256×256 FID of IntroVAE for LSUN Bedrooms but not for CELEBA-HQ (they provide it for 1024×1024 as 5.19, but with a more favorable train/test split). For other figures, the best-run results are shown. For GLOW, 256×256 LSUN image generation has not been demonstrated (only 128×128). FID is based on a 50k batch of generated samples compared to training samples. CELEBA-HQ Perceptual Path Length (PPL) was calculated with 100k samples, cropped to 128×128 ($\varepsilon = 10^{-4}$), LSUN PPL with 256×256 . Pre-trained models for PGGAN and GLOW ($T = 0.9$ resulted in best FID) were used with default settings provided by the authors. For all numbers, **smaller is better**.

	FID (CELEBA-HQ)	FID (LSUN)	PPL (CELEBA-HQ)	PPL (LSUN)
PGGAN	8.03	8.34	229.2	1080.1
IntroVAE	—	8.84	—	—
GLOW	68.93	—	219.6	—
PIONEER	39.17	18.07	155.2	779.5
Balanced PIONEER (ours)	25.25	17.89	146.2	678.4

the latent space as such, we measure the Perceptual Path Length [20] and show interpolations between reconstructions of input images as well as feature modifications. The success in the latter requires a well-structured latent space.

Since PIONEER networks are most useful in the domain of high-resolution images (128×128 and higher), we consider the CELEBA-HQ [19] and LSUN Bedrooms [42] datasets with images up to 1024×1024 and 256×256 resolution, respectively. CELEBA-HQ contains 30,000 images (where we use 27,000 / 3,000 split for training and testing images) while LSUN has its designated separate testing images. In Table 1, we compare the PIONEER, Balanced PIONEER, PGGAN and GLOW, with both datasets, in 256×256 .

We train progressively as in [13], allowing the last two pre-training phases (64×64 and 128×128) last 0.5–2 times longer than the earlier phases. The final phase is trained until convergence of the FID metric. We train both the CELEBA-HQ model and the LSUN Bedrooms model for 8 days on two Nvidia V100 GPUs, up to 25.5M samples (resolutions vary between epochs). In higher resolutions, the batch size is reduced. After the pre-training stages up to 64×64 , we switch on the margin ($m = 0.6$ for LSUN, $m = 0.2$ for CELEBA-HQ). From 256×256 onwards, $m = 0.4$. The fine-tuning is not necessary, but slightly improves the FID values. The overall scale of m can be determined empirically from the behavior of KL divergence values.

4.1. Ablation studies in varying the normalization

We carried out a sequence of ablation studies to quantify each single change in the normalization scheme. Fig. 3b



Figure 5. Generated samples from CELEBA-HQ and LSUN Bedrooms. IntroVAE from [16] (CELEBA-HQ downsampled), others uncurred.

focuses on the critical early training stages separately using PixelNorm (PN), Equalized Learning Rate (EQLR), PN+EQLR *etc.* in CELEBA-HQ up to 20M steps. Other hyper-parameters and resolution schedule remain the same. PN or EQLR alone or in combination will lead to poor FID even with KL margin. On the other hand, Spectral Normalization (SN) without KL margin will lead to divergence of the KL terms and training collapse. Only adding SN with the margin produces the optimal outcome. Also in PGGAN, replacing PN+EQLR with SN retained the FID results with CELEBA-HQ in 128×128 (see Supplement).

4.2. LSUN Bedrooms

For LSUN Bedrooms, we show randomly generated samples in Fig. 5b in 256×256 . In terms of visual comparison to, for example, [16], [19], or [11], it is hard to discern the differences between the models. However, comparing to 128×128 samples of [25] or [13] (which provide no 256×256 samples), we can already observe some mode collapse. We found that even the baseline PIONEER can indeed generate high-quality samples of LSUN Bedrooms at 256×256 ¹. The balanced PIONEER reaches better PPL and similar FID figures with 30% less training steps in the final stage than

¹This may be due to fixing the PyTorch bug 12671[39] which prevented the spectral normalization from working on multi-GPU setup.

PIONEER baseline (Table 1). In terms of FID scores, our Balanced PIONEER does not quite reach the state-of-the-art of [19] and [16]. However, the better performance of purely generative models such as [19] for random sampling is not directly comparable to that of encoder-based models.

4.3. CelebA-HQ

In [13], PIONEER was shown to reconstruct CELEBA images up to 128×128 , but only random samples at 256×256 level. We show reconstructions from both the baseline PIONEER and Balanced PIONEER at 256×256 . Balanced PIONEER has superior quality in both reconstructions (Fig. 2) and random sampling (Fig. 5a and Table 1). Reconstruction of input images is not interesting as such, but it serves as a proxy guarantee that the model at least handles the point in the latent space that represents the input image.

For identity conservation, the median L2 distance for 1k CELEBA-HQ samples in the embedding space of a DLib face recognition model improved from 0.758 of the PIONEER baseline to 0.712 of our model (29.1% closer to the target of 0.600 at which the recognizer is confident that the person is same; see the Supplement for details). LPIPS score for 10k CELEBA-HQ samples (cropped at face area to 128×128 for more precise measurement) improved from 0.223 of the baseline to 0.172 of our improved model (22.9% reduction).




Figure 6. Examples of feature manipulation in 256×256 resolution CELEBA-HQ test set images that the model has not seen before, for attributes that the model did not know about during training. Column 1: the input; Column 2: the reconstruction; Columns 3–6: various added feature vectors. Each 512-dimensional feature vector was extracted after training as a simple arithmetical difference between the latent vector of 32 to 64 training set examples that contain and do not contain the feature. The difference vector was then added to the latent code of each new input image and decoded back into a result image. There is no *a priori* reason to assume that this would result in a valid image in the first place, nor that the feature would be transformed. Note that there is no class information in the training set, which means that during training, the model has to figure out the existence of these features from scratch. Only a well-structured latent space can explain the result. Here, the feature intensity scale λ varies between $[-2.0, 2.0]$. For smoothly increasing the intensity, please see the Supplement.

4.4. Disentanglement measures

Finally, we evaluate whether the individual factors of variation in the training data are represented as disentangled directions in the latent space. We demonstrate this by showing smooth interpolations between the reconstructions of most input images (Fig. 1 and the Supplement) and attribute manipulation (Fig. 6) in CELEBA-HQ. We show superior degree of disentanglement with PPL, following [20]. To this end, we sample 100k pairs of short random latent vector segments, the ends of which we decode into pairs of images, cropped around mid-face. The scaled expectation of their LPIPS value is, then, the PPL.

5. Discussion and conclusions

In this paper, we introduced a streamlined PIONEER model variant, drawing from careful analysis of several normalization techniques and loss dynamics. Via state-of-the-art image manipulation capabilities, we also showed that AGE-based models can compete with GANs and VAEs.

We focused on face images to demonstrate the full range of capabilities of the model, even though it is in no way

specifically designed for face data. Unlike GANs limited to random sample manipulation, our model allows for encoding, reconstruction and manipulation of new real inputs. Unlike VAEs, we showed sharp reconstruction and attribute editing up to 256×256 resolution. The reconstructive capacity and latent space disentanglement are superior to the results shown so far with vanilla PIONEER and IntroVAE. Recent GANs and IntroVAE (but not GLOW) have shown better FID, but anything beyond random sample generation requires finer metrics. To evaluate latent space structure, arguably the core of representation learning, we showed unsupervised editing of face attributes (also see the Supplement), sharp interpolations between new inputs, and PPL metrics superior to the baselines on both datasets (Table 1).

The code for the experiments can be found at <https://github.com/AaltoVision/balanced-pioneer>.

Acknowledgements The authors acknowledge the Aalto Science-IT project, CSC – IT Center for Science, Finland, and GenMind Ltd for computational resources, and support by the Academy of Finland grants 308640, 324345, 277685, and 295081.

References

- [1] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning (ICML)*, pages 214–223, 2017.
- [2] M. Binkowski, D. J. Sutherland, M. Arbel, and A. Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations (ICLR)*, 2018.
- [3] A. Brock, J. Donahue, and K. Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations (ICLR)*, 2019.
- [4] A. Brock, T. Lim, J. M. Ritchie, and N. Weston. Neural photo editing with introspective adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- [5] T. Q. Chen, X. Li, R. B. Grosse, and D. K. Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2610–2620, 2018.
- [6] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo. StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8789–8797, 2018.
- [7] J. Donahue, P. Krähenbühl, and T. Darrell. Adversarial feature learning. In *International Conference on Learning Representations (ICLR)*, 2017.
- [8] V. Dumoulin, I. Belghazi, B. Poole, O. Mastropietro, A. Lamb, M. Arjovsky, and A. Courville. Adversarially learned inference. In *International Conference on Learning Representations (ICLR)*, 2017.
- [9] A. Geitgey. Face Recognition. https://github.com/ageitgey/face_recognition. GitHub repository.
- [10] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2672–2680, 2014.
- [11] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein GANs. In *Advances in Neural Information Processing Systems (NIPS)*, pages 5767–5777. 2017.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *International Conference on Computer Vision (ICCV)*, pages 1026–1034, 2015.
- [13] A. Heljakka, A. Solin, and J. Kannala. Pioneer networks: Progressively growing generative autoencoder. In *Asian Conference on Computer Vision (ACCV)*, pages 22–38, 2018.
- [14] A. Heljakka, A. Solin, and J. Kannala. Recursive chaining of reversible image-to-image translators for face aging. In *Proceedings of Advanced Concepts for Intelligent Vision Systems*, pages 309–320, 2018.
- [15] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6626–6637, 2017.
- [16] H. Huang, Z. Li, R. He, Z. Sun, and T. Tan. IntroVAE: Introspective variational autoencoders for photographic image synthesis. In *Neural Information Processing Systems (NeurIPS)*, pages 52–63, 2018.
- [17] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1125–1134, 2017.
- [18] D. Jimenez Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning (ICML)*, pages 1278–1286, 2014.
- [19] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations (ICLR)*, 2018.
- [20] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410, 2019.
- [21] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim. Learning to discover cross-domain relations with generative adversarial networks. In *International Conference on Machine Learning (ICML)*, pages 1857–1865, 2017.
- [22] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [23] D. Kingma and M. Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations (ICLR)*, 2014.
- [24] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

- [25] D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 10236–10245, 2018.
- [26] K. Kurach, M. Lucic, X. Zhai, M. Michalski, and S. Gelly. The GAN landscape: Losses, architectures, regularization, and normalization. In *Reproducibility in Machine Learning Workshop at ICML 2018*, 2018.
- [27] G. Lample, N. Zeghidour, N. Usunier, A. Bordes, L. Denoyer, and M. Ranzato. Fader networks: Manipulating images by sliding attributes. In *Advances in Neural Information Processing Systems (NIPS)*, pages 5969–5978, 2017.
- [28] A. Larsen, S. Kaae Sønderby, H. Larochelle, and O. Winther. Autoencoding beyond pixels using a learned similarity metric. In *International Conference on Machine Learning (ICML)*, pages 1558–1566, 2016.
- [29] F. Locatello, S. Bauer, M. Lucic, G. Rätsch, S. Gelly, and B. O. Schölkopf, B. Challenging common assumptions in the unsupervised learning of disentangled representations. pages 114–4124, 2019.
- [30] A. Makhzani, J. Shlens, N. Jaitly, and I. Goodfellow. Adversarial autoencoders. In *International Conference on Learning Representations (ICLR)*, 2016.
- [31] L. Mescheder, A. Geiger, and S. Nowozin. Which training methods for GANs do actually converge? In *International Conference on Machine Learning (ICML)*, pages 3481–3490, 2018.
- [32] L. Mescheder, S. Nowozin, and A. Geiger. Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks. In *International Conference on Machine Learning (ICML)*, pages 2391–2400, 2017.
- [33] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- [34] M. Rosca, B. Lakshminarayanan, D. Warde-Farley, and S. Mohamed. Variational approaches for auto-encoding generative adversarial networks. *arXiv preprint arXiv:1706.04987*, 2017.
- [35] K. Roth, A. Lucchi, S. Nowozin, and T. Hofmann. Stabilizing training of generative adversarial networks through regularization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2018–2028, 2017.
- [36] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2234–2242, 2016.
- [37] J. Tabor, S. Knop, P. Spurek, I. Podolak, M. Mazur, and S. Jastrzębski. Cramer–Wold AutoEncoder. *arXiv preprint arXiv:1805.09235*, 2018.
- [38] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.
- [39] Tongzhou Wang. Fix SpectralNorm with DataParallel 12671. <https://github.com/pytorch/pytorch/pull/12671>. GitHub repository.
- [40] D. Ulyanov, A. Vedaldi, and V. Lempitsky. It takes (only) two: Adversarial generator-encoder networks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI)*, pages 1250–1257, 2018.
- [41] Z. Yi, H. Zhang, P. Tan, and M. Gong. DualGAN: Unsupervised dual learning for image-to-image translation. In *International Conference on Computer Vision (ICCV)*, pages 2849–2857, 2017.
- [42] F. Yu, A. Seff, Y. Zhang, S. Song, T. Funkhouser, and J. Xiao. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- [43] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, and D. Metaxas. StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *International Conference on Computer Vision (ICCV)*, pages 5907–5915, 2017.
- [44] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018.
- [45] W. Zhou, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2008.
- [46] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *International Conference on Computer Vision (ICCV)*, pages 2223–2232, 2017.