



PAPER

OPEN ACCESS

RECEIVED
24 August 2023REVISED
6 December 2023ACCEPTED FOR PUBLICATION
14 December 2023PUBLISHED
4 January 2024

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Generating artificial displacement data of cracked specimen using physics-guided adversarial networks

David Melching^{*} , Erik Schultheis and Eric Breitbarth

German Aerospace Center (DLR), Institute of Materials Research, Linder Hoehe, 51147 Cologne, Germany

^{*} Author to whom any correspondence should be addressed.E-mail: David.Melching@dlr.de**Keywords:** physics-guided neural networks, generative adversarial networks, digital image correlation, fatigue crack growth

Abstract

Digital image correlation (DIC) has become a valuable tool to monitor and evaluate mechanical experiments of cracked specimen, but the automatic detection of cracks is often difficult due to inherent noise and artefacts. Machine learning models have been extremely successful in detecting crack paths and crack tips using DIC-measured, interpolated full-field displacements as input to a convolution-based segmentation model. Still, big data is needed to train such models. However, scientific data is often scarce as experiments are expensive and time-consuming. In this work, we present a method to directly generate large amounts of artificial displacement data of cracked specimen resembling real interpolated DIC displacements. The approach is based on generative adversarial networks (GANs). During training, the discriminator receives physical domain knowledge in the form of the derived von Mises equivalent strain. We show that this physics-guided approach leads to improved results in terms of visual quality of samples, sliced Wasserstein distance, and geometry score when compared to a classical unguided GAN approach.

1. Introduction

Fatigue crack growth (FCG) experiments are of significant importance to determine the lifetime and damage tolerance of critical engineering structures and components that are subjected to non-constant loads [1]. In recent years, digital image correlation (DIC) has been used to accompany and evaluate such mechanical experiments [2]. The DIC data serves as the basis for subsequent mechanical evaluation of fracture mechanical quantities like the stress intensity factors (SIFs) [3] and J -integral [4]. For this evaluation, the spatial location of the crack and especially the exact crack tip position is crucial. However, DIC data is subject to inherent noise and artefacts due to influences such as pattern quality, sensor noise, air movement, etc [5] which makes this information difficult to obtain. To improve upon these issues, extensive work was done optimizing the pattern quality [6] as well as the DIC algorithm in order to obtain reliable measurements in case of inferior patterns [7] or under special external conditions [8]. In the context of fracture mechanics, convolutional neural networks have been successfully applied to solve the crack detection problem fully automatically [9, 10]. These networks can deal with DIC noise and take the interpolated DIC-measured displacement fields as input to predict the crack paths and tips. However, for these powerful data-driven models to work reliably, they need a diverse set of training data. At the same time, experimental training data is scarce, since experiments are expensive, and manual labelling is extremely tedious and time-consuming. To address this issue, Strohmann *et al* [9] added artificial training data in the form of finite element (FE) simulations. Nevertheless, simulations are idealized and lack the characteristic DIC noise.

Classically, synthetic DIC data can be generated by first creating an artificial speckle pattern on a digital image of the desired specimen. Then, an FE simulation is used to virtually deform the image. Finally, the deformed (speckle) image can be evaluated using a DIC algorithm (see, e.g. [11] and more recently [12]). This synthetic DIC data can for instance be used to assess systematic errors arising from different DIC techniques [11] or to understand calibration uncertainty [13]. In contrast to these approaches, our goal is to bypass this multi-step procedure altogether and directly generate large amounts of artificial interpolated

DIC-like training data in a fast, simple, data-centric manner using machine-learned generative models. For instance, this interpolated DIC data can be used to improve the training of automatic crack detection models (see [9, 10]) by delivering an advanced data augmentation method.

In the field of data-driven modelling, generative adversarial networks (GANs) have proven to be powerful data generators. GANs are a generative, unsupervised approach to machine learning based on deep neural networks trained using an adversarial process. Deep convolutional GANs (DC-GANs) produced state-of-the-art results in many generative and translative computer vision tasks such as image generation and style transfer [14–16]. However, training GANs typically requires large amounts of data, which are often not available in the scientific domain. For example, it is not possible to mechanically test an entire aircraft to generate training data. But without sufficient data, deep learning models are often unreliable and poorly generalize to domains not covered by the training data. To overcome this problem, efforts have been made to integrate fundamental physical laws and domain knowledge into machine learning models. Karpatne *et al* [17] describe theory-guided data science as an emerging paradigm aiming to use scientific knowledge to improve the effectiveness of data science models. They presented several approaches for integrating domain knowledge in data science. Daw *et al* [18] proposed a physics-guided neural network by adding a physics-based term to the loss function of a neural network to encourage physically consistent results. Karniadakis *et al* [19] coined the term physics-informed neural networks—a deep learning framework which enables the combination of data-driven and mathematical models described by partial differential equations. Yang *et al* [20] combined this approach with GANs by adding a physics-based loss to the generator and, recently, Daw *et al* [21] introduced a physics-informed discriminator GAN, which physically supervises the discriminator instead.

In this work, we generate artificial displacement data of cracked specimen using GANs. Our framework is based on the classical DC-GAN architecture. We incorporate mechanical knowledge by using a physics-guided discriminator. In addition to the generated displacement data, this discriminator receives the equivalent strain according to Mises [22] derived from the real or fake (generated) displacements as an additional input feature, see section 2 and figure 1. This mechanically motivated physical feature guides the adversarial training process and leads to physically consistent generated data. The generated data samples can be used to increase data variation of given training datasets consisting of interpolated DIC displacements. Although this synthetic data is not labelled, it has the potential to improve supervised machine learning tasks, e.g. by using unsupervised pre-training [23] or label propagation [24]. In general, our method can be applied to generate interpolated DIC-like displacement fields of cracked specimen. In this paper, we focus on FCG experiments of a middle tension (MT) specimen manufactured from an aluminium-based alloy (see section 2.2) and train several GANs. To demonstrate the merits of the physics-guided method, we compare the results of the physics-guided GAN to a classical GAN approach in terms of visual quality of generated samples (see section 3.2) and distance of fake data distributions to the real training data. The latter is quantified using the sliced Wasserstein distance (SWD) [25] and the geometry score (GS) [26] (see section 2.4). We show that the physics-guided approach accelerates the training and leads to physically more consistent results.

2. Methodology

GANs are generative machine learning models learned using an adversarial training process [27]. In this framework, two neural networks—the generator G and the discriminator D —contest against each other in a zero-sum game. Given a training dataset characterized by a distribution p_{data} , the generator aims to produce new data following p_{data} while the task of the discriminator is to distinguish generated data samples from actual training data samples.

Given a noise vector z sampled from a prior, e.g. the standard normal distribution, the generator outputs data samples $G(z)$, called *fake samples*, trying to follow the training data distribution p_{data} . Given a real or fake sample, the discriminator is supposed to decide whether it is real or fake by predicting the probability of it belonging to the training dataset.

Both models G and D are trained simultaneously contesting against each other in a two-player zero-sum minimax game with the value function $V(G, D)$:

$$\min_G \max_D V(G, D) := \mathbb{E}_{x \sim p_{\text{data}}} [\log(D(x))] + \mathbb{E}_z [\log(1 - D(G(z)))]. \quad (1)$$

This means D is trained to minimize the discriminator loss

$$L_D = -\log(D(x)) - \log(1 - D(G(z))), \quad (2)$$

whereas G is trained to minimize the generator loss

$$L_G = \log(1 - D(G(z))). \quad (3)$$

As the discriminator gets better at identifying fake samples $G(z)$, the generator has to improve on generating samples which are more similar to the real training samples $x \sim p_{\text{data}}$. We refer to [27] for further details of the training algorithm.

2.1. DIC

DIC is a contact-less, optical method to obtain full-field displacements and strains. It is widely applied in science and engineering to quantify deformation processes. In experimental mechanics, it is used to monitor and evaluate FCG experiments [28] by determining fracture mechanical parameters like SIFs [3] or the J-integral [4]. Essentially, DIC measurements are based on the comparison of a current image with a reference image using tracking and image registration techniques. The cross correlation method requires a random speckle pattern on the sample surface. Various external and internal influences such as illumination, air movement, vibrations, facet size and spacing, camera settings, sensor noise, pattern quality, etc lead to inherent noise in the DIC data. Our goal is to generate artificial interpolated DIC-like displacement data using GANs. Since this data incorporates characteristic DIC noise, it can subsequently be used to improve the training of machine learning models such as crack detection [9, 10].

2.2. Training data

To create a dataset for the training of our GANs, we use planar displacement fields $u = (u_x, u_y)$ obtained during FCG experiments of the aluminium alloy AA2024-T3 using a commercial GOM Aramis 12M 3D-DIC system. Details on the general experimental setup can be found in [9]. For the dataset, we use *one* FCG experiment performed on an MT specimen (width $w = 160$ mm, thickness $t = 2$ mm). While the specimen is clamped at the bottom, a load is applied from the top with a maximal force of $F_{\text{max}} = 15$ kN (corresponding to $\sigma_{\text{max}} = 47$ MPa) and ratio $R = 0.3$ with 20 load cycles per second. Every 0.5 mm of crack growth (measured by direct current potential drop), three images (at minimal load $F_{\text{min}} = R \cdot F_{\text{max}}$, mean load $F_{\text{min}} + (F_{\text{max}} - F_{\text{min}})/2$ and maximum load F_{max}) were acquired. From the resulting DIC dataset, we take the planar displacements u_x and u_y of the specimen and linearly interpolate them from an area of $70 \times 70 \text{ mm}^2$ of the right-hand side of the specimen on an equidistant 256×256 pixel grid. This procedure results in 838 data samples of shape $2 \times 256 \times 256$, where the first dimension stands for the x - and y -displacements. Each of the two channels is normalized to $[-1, 1]$ by the min-max-scaling and shift

$$u_{\text{scaled}} = 2 \cdot \frac{u - u_{\text{min}}}{u_{\text{max}} - u_{\text{min}}} - 1 \quad (4)$$

such that the minimum and maximum are mapped to -1 and 1 , respectively.

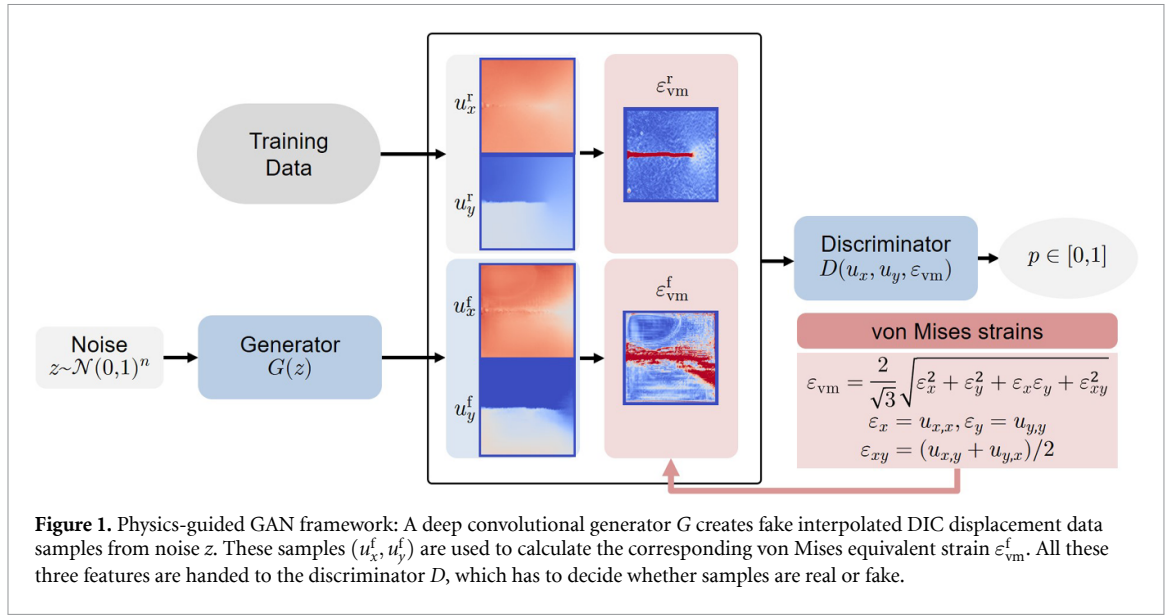
2.3. Physics-guided GAN

We aim to generate artificial interpolated DIC displacement data using DC-GANs. For this, we mainly follow the architectural guidelines from [29]. However, in order to reduce checkerboard artefacts, we choose nearest-neighbour upsampling instead of transposed convolutions in the generator [30]. We remark that GANs cannot be expected to generalize beyond the training data. The reason for this is that the generator learns to produce fake samples that approximate the distribution of the training data. To cover a different type of experiment or material, the training data set must be extended.

Generator. The input of the generator network is a n -dimensional vector z randomly sampled from a standard normal distribution. For definiteness, we choose $n = 5$ throughout all our training experiments. First, the random vector z passes a fully-connected layer with $8 \times 8 \times 512 = 32768$ neurons, batch-normalization [31], and rectified linear unit (ReLU) activation [32]. The output of this layer is then reshaped into 512 features of size 8×8 . After that, these features are successively doubled in size using the base block (upsampling \rightarrow batch normalization \rightarrow ReLU \rightarrow convolution) four times. The final block ends with a tanh activation instead of ReLU. Therefore, in accordance with the training data, the generator outputs *fake* samples with pixel values between -1 and 1 .

Discriminator. For the discriminator, we implemented the following two approaches:

1. **Classical:** The discriminator gets real and fake pairs of interpolated x - and y -displacement fields (u_x, u_y) and predicts a (pseudo-)probability of the sample being real. We refer to this approach as *classical* GAN.
2. **Physics-guided:** In addition to the interpolated displacement fields, the corresponding von Mises equivalent strain ε_{vm} is calculated based on the generated and real interpolated displacement fields and



the discriminator gets the triple $(u_x, u_y, \epsilon_{vm})$ as input in order to decide whether it is fake or real. For small strains, the von Mises equivalent strain is defined as the scalar quantity

$$\epsilon_{vm} = \sqrt{\frac{2}{3} \epsilon_{dev} : \epsilon_{dev}}, \quad \text{where } \epsilon_{dev} = \epsilon - \frac{1}{3} \text{tr}(\epsilon) \quad (5)$$

denotes the deviatoric part of the three-dimensional strain tensor

$$\epsilon = \begin{pmatrix} \epsilon_{xx} & \epsilon_{xy} & \epsilon_{xz} \\ \epsilon_{xy} & \epsilon_{yy} & \epsilon_{yz} \\ \epsilon_{xz} & \epsilon_{yz} & \epsilon_{zz} \end{pmatrix}. \quad (6)$$

In case of plane stress, $\epsilon_{xz} = \epsilon_{yz} = 0$ and $\epsilon_{zz} = -\nu(\epsilon_{xx} + \epsilon_{yy})$. Assuming volume constancy with a Poisson ratio of $\nu = 1/2$, formula (5) simplifies to

$$\epsilon_{vm} = \frac{2}{\sqrt{3}} \sqrt{\epsilon_{xx}^2 + \epsilon_{yy}^2 + \epsilon_{xy}^2 + \epsilon_{xx}\epsilon_{yy}}. \quad (7)$$

We use formula (7) for the physics-guided discriminator. Therefore, we numerically approximate the strains using finite differences, e.g.

$$\epsilon_{xy}(x, y) = \frac{\partial}{\partial y} u_x(x, y) = \frac{u_x(x, y+h) - u_x(x, y)}{h} + \mathcal{O}(h). \quad (8)$$

To guarantee differentiability, the square-root function is smoothed by using $\sqrt{\cdot + \delta}$ with $\delta \ll 1$. We refer to this approach as *physics-guided* GAN (see figure 1). The discriminator, which drives the training of the generator, has additional physical information, namely the equivalent strain, that the generator can only influence indirectly by producing physically consistent displacement fields. Certainly, other quantities, which can be derived from the displacement fields, such as strains ϵ_{xx} , ϵ_{xy} , or ϵ_{yy} can be used to guide the discriminator. However, we decide on using the equivalent strains, since the crack path and crack tip field is well-visible in them.

In both cases, we choose the same model architecture for the discriminator. The input of size $2 \times 256 \times 256$, or $3 \times 256 \times 256$ in case of the physics-guided discriminator, is successively downsampled to the size $1 \times 32 \times 32$ using three blocks of strided convolutions, batch normalization, and LeakyReLU activation [33], where $\text{LeakyReLU}(t) = \max(\alpha t, t)$ with $\alpha = 0.2$. The extracted features are then flattened and pass the last fully-connected layer with one output neuron and sigmoid activation. The output is a number between 0 and 1 and is interpreted as the probability of the sample being real.

2.4. Evaluation of GANs

In classical supervised learning, a model is trained by minimizing a specific loss (e.g. mean squared error), which quantitatively compares model predictions with the expected target. After training, models can be evaluated and compared by calculating the loss (and accuracy) for independent labelled test data. GAN generators, however, are trained in an adversarial fashion using a second model (the discriminator) to classify generated data as real or fake. Both models are trained simultaneously to maintain an equilibrium. Therefore, there is no natural objective measure to evaluate GANs, quantitatively. Instead they are evaluated by assessing the quality and variation of generated data. This is typically achieved by visual inspection of generated samples or by calculating the inception score (IS) [34] and Fréchet inception distance (FID) [35]. However, in case of interpolated DIC data, several domain experts would be needed to objectively grade the visual quality of generated samples. Moreover, quantitative metrics like IS or FID can only be employed for natural images since they use image classification networks like Inception [36], which are pre-trained on ImageNet [37]. Therefore, in addition to a visual examination of generated samples in section 3.2, we use metrics which are independent of the data type and do not use any pre-trained models. More precisely, we use the following two metrics:

Sliced Wasserstein distance. In mathematics, the Wasserstein distance is a natural distance function between two distributions. Intuitively, it can be viewed as the minimal cost of transforming one of the distributions into the other. In case of image-like datasets $X = \{X_n\}_{n=1,\dots,N}$ and $Y = \{Y_n\}_{n=1,\dots,N}$ with same number of samples N and image sizes $c \times h \times w$, where c is the number of channels and h and w denote the height and width of images, respectively, the (quadratic) Wasserstein distance is given by,

$$W(X, Y)^2 = \min_{\pi} \sum_{i,j,k,n} |X_n(i, j, k) - Y_{\pi(n)}(i, j, k)|^2, \quad (9)$$

where the minimum is taken over all permutations π of the set $\{1, \dots, N\}$ [25]. Due to the high dimensionality of images and the large number of samples, the exact computation of the Wasserstein distance is computationally infeasible. This is because the number of permutations scales exponentially with the number of samples N . Therefore, instead of (9), we use the SWD introduced in [25] as an approximation, which is amenable for efficient numerical computation. The main idea of slicing is to map the high dimensional image data from $\mathbb{R}^{c \times h \times w}$ onto one-dimensional slices. On these slices, the Wasserstein distance can be calculated in loglinear time by using the ordered structure of one-dimensional Euclidean space. The SWD is defined as,

$$\tilde{W}(X, Y)^2 = \int_{\theta \in \Omega} \min_{\pi_{\theta}} \sum_{n=1}^N |\langle X_n - Y_{\pi_{\theta}(n)}, \theta \rangle|^2 d\theta, \quad (10)$$

where $\Omega = \{\theta \in \mathbb{R}^{c \times h \times w} : \|\theta\| = 1\}$ denotes the unit sphere. We refer to [14, 25] and section 3.3 below for further details.

Geometry score. Introduced by Khrulkov and Oseledets [26], the GS allows to quantify the performance of GANs trained on datasets of arbitrary nature. It measures the similarity between the real dataset X_{real} and a generated one X_{fake} by comparing topological properties of the underlying low-dimensional manifolds [38]. The detailed quantitative characterization of the underlying manifold of a given dataset X is usually very hard. The core idea of [26] is to choose random subsets $L \subset X$ called *landmarks* and to build a family of *simplicial complexes*, parametrized by a non-negative, time-like *persistence parameter* α . For small α , the complexes consist of a disjoint union of points. Increasing α adds more and more simplices finally leading to one single connected blob. For each value of α , topological properties of the corresponding simplicial complex, namely the number of *one-dimensional holes* in terms of homology, $\beta_1(\alpha)$, are calculated (see, e.g. [39]). From this, the authors propose to compute relative living times (RLTs) for every number of holes that was observed [26]. For each non-negative number i , the RLT is the amount of the time when exactly i holes were present relative to the overall time α_{max} after which everything is connected. More precisely,

$$\text{RLT}(i, X, L) = \frac{\mu(\{\alpha \in [0, \alpha_{\text{max}}] : \beta_1(\alpha) = i\})}{\alpha_{\text{max}}}, \quad (11)$$

where μ denotes the standard Lebesgue measure. Since the RLTs depend on the choice of landmarks L , we choose a collection of n random sets of landmarks L_j and define the mean relative living times (MRLTs) as

$$\text{MRLT}(i, X) = \frac{1}{n} \sum_{j=1}^n \text{RLT}(i, X, L_j). \quad (12)$$

The MRLT is a discrete probability distribution over the non-negative integers. It can be interpreted as the probability of the manifold having exactly i one-dimensional holes (on average). The L^2 -distance between the MRLT distributions of X_{real} and X_{fake} defines a measure of topological similarity between the real dataset and the generated one, called GS:

$$\text{GS}(X_{\text{fake}}, X_{\text{real}}) = \sum_{i=0}^{i_{\text{max}}-1} |\text{MRLT}(i, X_{\text{fake}}) - \text{MRLT}(i, X_{\text{real}})|^2, \quad (13)$$

where i_{max} is an upper bound on the number of holes. We refer to [26] for further theoretical details and to section 3.4 for the choice of hyperparameters and results in our case.

3. Results and discussion

In order to demonstrate the effectiveness of the method and to compare the classical with the physics-guided discriminator approach, we trained 10 randomly initialized classical and physics-guided GANs each for 100 epochs. Moreover, we trained two classical and physics-guided GANs each for 1000 epochs in order to compare both architectures after long training runs. The training setup is described in section 3.1 below. The trained models are evaluated qualitatively and quantitatively by using the following criteria:

- Visual inspection of randomly generated samples (section 3.2)
- Sliced Wasserstein distances (section 3.3)
- Geometry scores (section 3.4)

A summary of the results can be seen in table 1. In short, the physics-guided GAN approach leads to visually better results after 100 epochs and overall to measurably better results. For a detailed discussion, we refer to the sections below.

3.1. Training procedure

Before training, the filters of the convolution layers in both generator and discriminator network are initialized randomly from a normal distribution with zero mean and a standard deviation 0.02. In contrast, the weights of the batch normalization layers are initialized from mean 1 and standard deviation of 0.02, whereas the biases are initialized with zeros.

For training, we choose the Adam optimizer [40] with a learning rate of 0.002, momentum parameters of $\beta_1 = 0.5$, $\beta_2 = 0.999$, and a batch size of 8. We noticed that occasionally models suffer from mode collapse during training. This means that the generator always outputs the same or visibly similar fake data samples and stops to learn. This problem is well-known and still part of active research. Popular strategies to overcome convergence issues of GANs regularize or perturb the discriminator [41, 42] or by using a more sophisticated loss function [43]. In our case, if mode collapse happened, we restarted the training and discarded the collapsed model. All neural networks and training loops were implemented using PyTorch [44]. The hardware for the training was an NVIDIA RTX8000 graphics card.

3.2. Visual evaluation

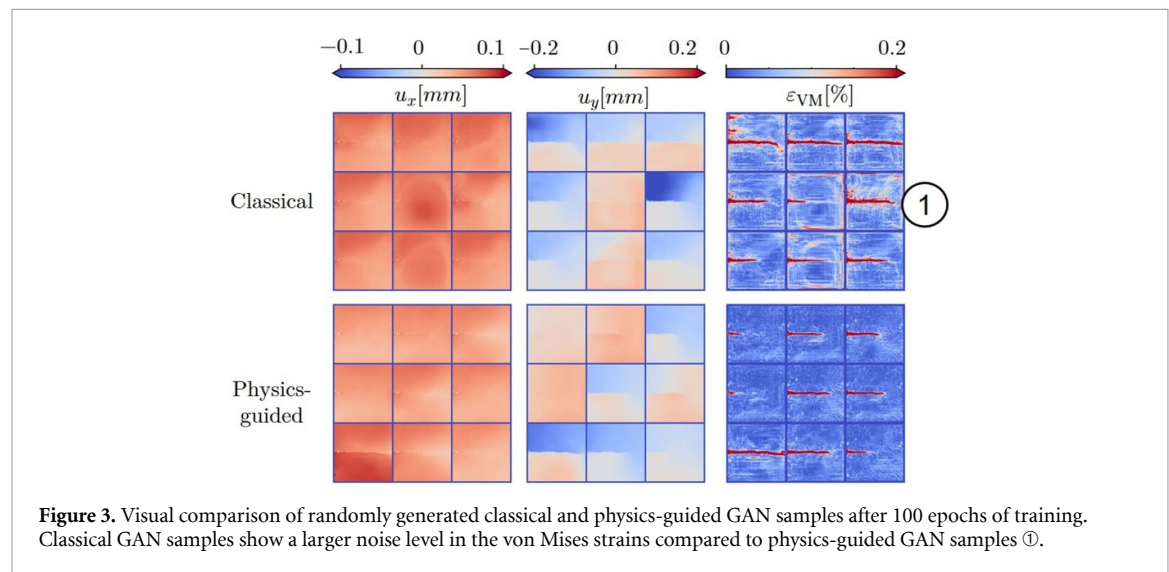
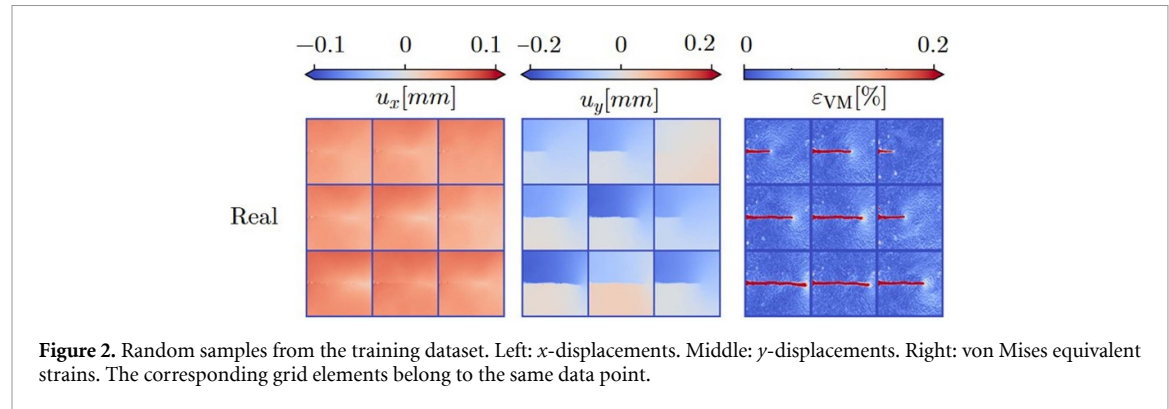
We begin with a visual inspection of the generated data and compare real training data to representative samples generated by the classical GAN and the physics-guided GAN. We refer to fake samples generated by the classical or physics-guided GAN generators as classical or physics-guided GAN samples, respectively.

Figure 2 shows real interpolated DIC data samples obtained during FCG experiments as described in section 2.2. The figure contains planar displacements and von Mises equivalent strains of nine randomly selected data samples. Images belong together in the sense that the x -displacement of the first sample is located at the top left of the left column. The corresponding y -displacement is located at the same position in the middle column, and the corresponding calculated equivalent strain is located at the same position in the right column. Here, the crack path as well as the characteristic crack tip field is clearly visible.

In figure 3, we see random fake samples after 100 epochs of GAN training. We can often identify the initial crack on the left edge and the crack path. Whereas most generated displacements are visually close to real displacements, significant differences are revealed in the von Mises strains, which are calculated afterwards. Especially classical GAN samples contain inconsistencies between x - and y -displacements and visual artefacts. This leads to large-scale vortexes and small-scale noise in the von Mises strains ①. Although far from being perfect, physics-guided GAN samples contain significantly less of these artefacts and inconsistencies and visually capture the inherent noise of the DIC system much better than classical GAN

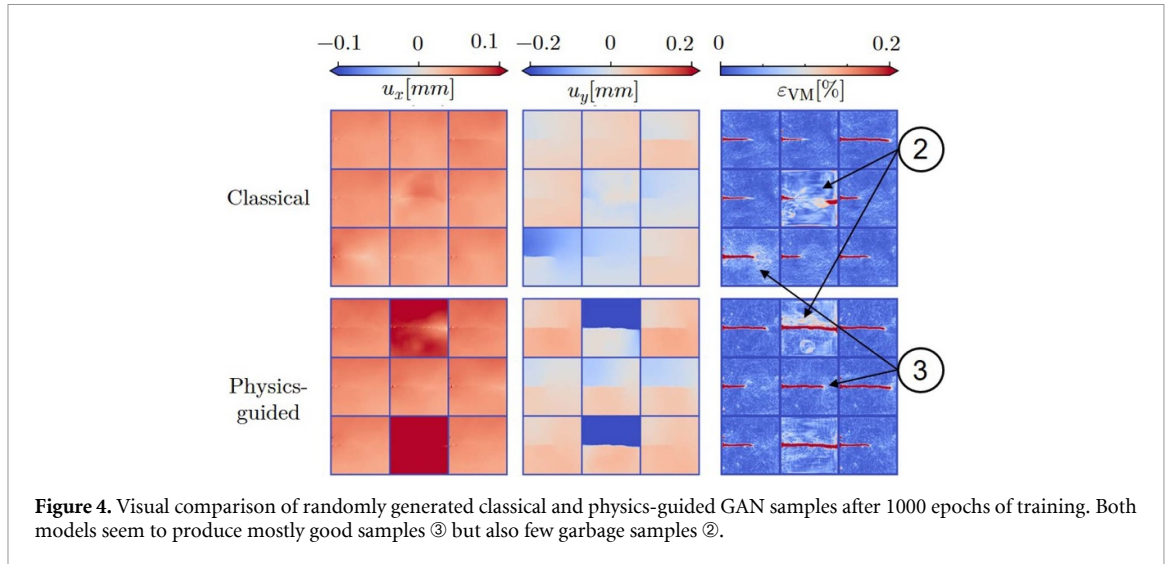
Table 1. Subjective visual quality, calculated geometry score (GS), and sliced Wasserstein distance (SWD) (lower is better) for different GAN model architectures and training lengths. (*) apart from *garbage* samples (see figure 4). The best results are highlighted in bold.

Model	Epoch	Visual quality	GS $\times 10^3$	SWD $\times 10^3$
Classical	100	Low	183.49	151.70
Physics-guided	100	Medium	157.58	82.96
Classical	1000	High*	23.83	83.19
Physics-guided	1000	High*	1.93	61.61



samples. Nevertheless, most fake samples are still visually distinguishable from real samples. In order to make sure the models are fully converged, we also performed some longer training runs.

Figure 4 shows random fake samples after 1000 epochs. At this stage, the models are well converged and the visual difference between classical and physics-guided GAN samples has mainly disappeared. In general, the fake samples of both models show much better visual quality and less artefacts and inconsistencies compared to fake samples of generators trained for only 100 epochs ③. However, few samples suffer from severe inconsistencies and are qualitatively inferior ②. We refer to these failures as *garbage* samples. Garbage samples may arise from difficulties in the training process of GANs and problems of non-convergence like mode collapse, which is an open research problem [34]. Although we do not observe mode collapse for the results shown here, the occurrence of garbage samples is a sign of local non-convergence, i.e. some noise inputs are mapped to garbage samples. Since the models are converged after 1000 epochs w.r.t. the difference in output between two epochs, the garbage samples seem to originate from intrinsic difficulties in the training process and are not related to the number of training epochs. Apart from these outliers, the vast majority of samples (of both models) are visually indistinguishable from real samples. Nevertheless, domain experts may notice that the characteristic crack tip field still seems unphysical in the fake samples especially when compared to real samples with long cracks.



3.3. SWDs

For a thorough comparison of GANs, one needs to inspect a large number of fake samples. Doing this manually, would be very tedious and subjective. Instead, one should compare the results using meaningful, quantitative metrics.

For this, we follow [14] and calculate the SWDs introduced in section 2.4 between fake data samples and real data samples on various scales. These scales are introduced by building a five-level Laplacian pyramid [45] with resolutions 16×16 , 32×32 , 64×64 , 128×128 , 256×256 . Each pyramid level corresponds to a specific spatial resolution. For each level, we compute the SWD between the training dataset and a generated fake dataset of the same size. More precisely, the SWDs are calculated between datasets of random 7×7 patches of the pyramid samples. The patches are pre-processed by normalizing each channel (i.e. x and y displacement) to mean 0 and standard deviation 1. To reduce uncertainty, we average the SWDs of ten runs with randomly sampled fake data. Since there are less unique patches for low resolutions, we adapt the number of random patches depending on the pyramid level. For the five resolutions, 16×16 , 32×32 , 64×64 , 128×128 , and 256×256 , we use 128, 256, 512, 1024, 2048 patches, respectively. The integral in equation (10) is approximated by choosing 512 random slices and averaging the results. We implemented a GPU-enabled version of the code from [14] using the PyTorch [44] framework.

At least intuitively, a small SWD shows that the fake and real samples are similar. At low resolution (e.g. 16×16) only large-scale features like the crack length are visible and a small SWD would indicate that the variation of crack lengths in the fake dataset is similar to the training dataset. At high resolution (e.g. 256×256) very fine-grained structures like the inherent DIC noise is encoded in the patches.

Figure 5 shows the calculated SWDs of classical and physics-guided GAN samples after 100 epochs of training. In order to estimate uncertainty, we trained ten randomly initialized models each with the classical and the physics-guided GAN architecture. The main observation is that for all resolutions, physics-guided samples are closer to the training data than classical GAN samples. This indicates that physics-guided GAN samples are better in quality and variation. Especially for the high resolution 256×256 , the SWDs show a large gap and confirm our visual observation of artefacts and unphysical noise as seen in the classical GAN samples in figure 3. Nevertheless, the results can be significantly different for each trained generator. This fact is reflected in the large error bars of the SWDs.

The results after 1000 training epochs are displayed in figure 6. Here, we used two training runs for each GAN architecture. As expected, the distances are all smaller than after 100 epochs. In contrast to the results after 100 epochs (cf figure 5), both GAN architectures are closer together. However, the physics-guided GAN samples have significantly smaller SWDs for the fine resolution 256×256 and the low resolutions 16×16 and 32×32 . This suggests that after 1000 epochs of training the physics-guided samples are still closer to the real samples in terms of quality and variation. Nevertheless, the few garbage samples seen in figure 4 ② could influence the SWDs especially at larger pyramid levels, e.g. 256×256 .

3.4. GSs

To compare the GS introduced in section 2.4 of different trained GANs, we generated fake datasets with the same number of samples $N = 838$ as the training dataset. To calculate the MRLTs of the real and fake datasets, we mainly follow the recommendations in [26]. We set $i_{\max} = 100$ and use $n = 1000$ random landmarks. The




Figure 5. Comparison of SWDs between classical GAN (left) and physics-guided GAN (right) trained for 100 epochs. The boxplot intervals range from the minimal to the maximal SWDs. The box includes ranges from the 25% to the 75% quantile and shows the median.




Figure 6. Comparison of SWDs between classical GAN (left) and physics-guided GAN (right) trained for 1000 epochs. The boxplot intervals range from the minimal to the maximal SWDs. The box includes ranges from the 25% to the 75% quantile and shows the median.

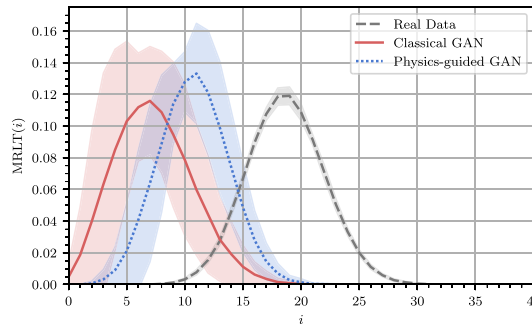


Figure 7. Comparison of MRLT distributions between the real dataset and fake datasets generated by classical and physics-guided GAN after 100 epochs of training.

number of samples in each landmark is 64. The maximal persistence time α_{\max} is proportional to the maximal pairwise Euclidean distance between samples in each landmark, i.e. for $j = 1, \dots, n$:

$$\alpha_{\max}^j = \gamma \max(\text{dist}(L_j, L_j)), \quad \gamma = \frac{1}{128} / \frac{N}{5000}. \quad (14)$$

We used the implementation from [26] to calculate the MRLTs.

Figure 7 shows the distributions of MRLTs after 100 epochs of training. The error band originates from the uncertainty induced by the random landmarks and, even more so, from the ten different models trained for each GAN architecture. This results in large variations of calculated MRLTs. Nevertheless, on average the physics-guided GAN distribution is closer to the MRLTs of the real data distribution than the classical GAN distribution. This observation is quantitatively reflected in a smaller mean GS of the physics-guided models (see table 1). However, both fake data distributions are still far away from the real data distribution and the GSs are large.




Figure 8. Comparison of MRLT distributions between the real dataset and fake datasets generated by classical and physics-guided GAN after 1000 epochs of training.

In figure 8, we see the MRLT distributions after 1000 epochs of training. Both GAN results are much closer to the real data than after 100 epochs and the physics-guided GAN MRLTs almost coincide with the real data MRLTs. This accordance is shown in the calculated GSs in table 1 as well.

4. Conclusion

We introduced a machine learning framework to generate artificial full-field displacements of cracked specimen by learning the underlying data distribution from a sufficiently large DIC dataset. The training data was obtained during FCG experiments of the aluminium alloy AA2024-T3. In contrast to FE simulations, our method is able to produce large amounts of interpolated DIC-like displacement data in a fast and easy way but is limited in the sense that boundary conditions and crack configurations cannot be controlled.

Our approach is based on DC-GANs. The main novelty compared to the classical DC-GAN framework is a physics-guided discriminator. This discriminator, in addition to the generated x - and y -displacement fields, gets also the derived von Mises equivalent strain as input. This enables the discriminator to detect physical inconsistencies in the generated fake samples more easily, thus enhancing the training process.

In order to evaluate trained generator models on an objective basis, we used two quantitative metrics. First, the SWD between real and fake samples and, secondly, the GS approximating the topological distance between a generated data manifold and the training data manifold.

We observed superior performance of the physics-guided GAN compared to the classical GAN approach. This result was observed by visual evaluation of generated samples and confirmed by lower SWDs and GSs of the physics-guided models. Both, SWD and GS, proved themselves to be valuable evaluation metrics. They are useful to identify mode collapse and to select the best trained models. Nevertheless, it is important to note that there is no natural metric to evaluate the performance of GANs. In the absence of powerful pre-trained models like Inception for DIC-like data, we had to stick to GAN metrics that are independent of these benchmark models. Our findings support the claim that hybrid models, which combine data-driven methods with physical domain knowledge, can lead to more powerful models and faster training.

The visual inspection revealed a varying sample quality. Especially the converged models after 1000 epochs of training, apart from mostly good samples, produce few garbage samples. Although the number of these garbage samples is model-dependent, we were not able to avoid their occurrence completely. Moreover, we still face the issue of (local) non-convergence and mode collapse. To overcome these issues, one could try to stabilize training using suitable regularization techniques [42, 46].

The main open problem concerns the control of boundary conditions like the crack path and external force. In contrast to FE-based data generation, with our approach it is not possible to control them. This challenge could be tackled by using a conditional GAN framework [47] and is part of current research.

Data availability statement

The code will be publicly available on Github (<https://github.com/dlr-wf>).

The data that support the findings of this study will be openly available following an embargo at the following URL/DOI: <https://doi.org/10.5281/zenodo.7737880>.

Acknowledgments

We acknowledge the financial support of the DLR-Directorate Aeronautics.

Competing interests

The authors declare no competing financial or non-financial interests.

Author contributions

D M and E S conceived the physics-guided GAN. E S implemented the neural network architectures, training algorithms, and evaluation metrics. All authors discussed, analysed, and interpreted the results and wrote the manuscript.

ORCID iDs

David Melching  <https://orcid.org/0000-0001-5111-6511>

Erik Schultheis  <https://orcid.org/0009-0007-4728-7124>

Eric Breitbarth  <https://orcid.org/0000-0002-3479-9143>

References

- [1] Tavares S M O and de Castro P M S 2017 T An overview of fatigue in aircraft structures *Fatigue Fract. Eng. Mater. Struct.* **40** 1510–29
- [2] Sutton M A and Hild F 2015 Recent advances and perspectives in digital image correlation *Exp. Mech.* **55** 1–8
- [3] Roux S, Réthoré J and Hild F 2009 Digital image correlation and fracture: an advanced technique for estimating stress intensity factors of 2D and 3D cracks *J. Phys. D: Appl. Phys.* **42** 214004
- [4] Becker T, Mostafavi M, Tait R and Marrow T 2012 An approach to calculate the J-integral by digital image correlation displacement field measurement *Fatigue Fract. Eng. Mater. Struct.* **35** 971–84
- [5] Zhao J, Sang Y and Duan F 2019 The state of the art of two-dimensional digital image correlation computational method *Eng. Rep.* **1** e12038
- [6] Chen Z, Shao X, Xu X and He X 2018 Optimized digital speckle patterns for digital image correlation by consideration of both accuracy and efficiency *Appl. Opt.* **57** 884
- [7] Su Z, Lu L, He X, Yang F and Zhang D 2020 Recursive-iterative digital image correlation based on salient features *Opt. Eng., Bellingham* **59** 034111
- [8] Su Z, Pan J, Lu L, Dai M, He X and Zhang D 2021 Refractive three-dimensional reconstruction for underwater stereo digital image correlation *Opt. Express* **29** 12131
- [9] Strohmann T, Starostin-Penner D, Breitbarth E and Requena G 2021 Automatic detection of fatigue crack paths using digital image correlation and convolutional neural networks *Fatigue Fract. Eng. Mater. Struct.* **44** 1336–48
- [10] Melching D, Strohmann T, Requena G and Breitbarth E 2022 Explainable machine learning for precise fatigue crack tip detection *Sci. Rep.* **12** 9513
- [11] Lava P, Cooreman S, Coppieters S, De Strycker M and Debruyne D 2009 Assessment of measuring errors in DIC using deformation fields generated by plastic FEA *Opt. Lasers Eng.* **47** 747–53
- [12] Rohe D and Jones E 2022 Generation of synthetic digital image correlation images using the open-source blender software *Exp. Tech.* **46** 615–31
- [13] Balcaen R, Wittevrongel L, Reu P L, Lava P and Debruyne D 2017 Stereo-DIC calibration and speckle image generator based on FE formulations *Exp. Tech.* **57** 703–18
- [14] Karras T, Aila T, Laine S and Lehtinen J 2018 Progressive growing of GANs for improved quality, stability and variation *Int. Conf. on Learning Representations*
- [15] Karras T, Laine S, Aittala M, Hellsten J, Lehtinen J and Aila T 2020 Analyzing and improving the image quality of StyleGAN *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*
- [16] Gatys L A, Ecker A S and Bethge M 2016 Image style transfer using convolutional neural networks *2016 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* pp 2414–23
- [17] Karpadne A, Atluri G, Faghmous J H, Steinbach M, Banerjee A, Ganguly A, Shekhar S, Samatova N and Kumar V 2017 Theory-guided data science: a new paradigm for scientific discovery from data *IEEE Trans. Knowl. Data Eng.* **29** 2318–31
- [18] Daw A, Karpadne A, Watkins W, Read J and Kumar V 2017 Physics-guided neural networks (PGNN): an application in lake temperature modeling (arXiv:1710.11431)
- [19] Karniadakis G E, Kevrekidis I G, Lu L, Perdikaris P, Wang S and Yang L 2021 Physics-informed machine learning *Nat. Rev. Phys.* **3** 422–40
- [20] Yang Y and Perdikaris P 2019 Adversarial uncertainty quantification in physics-informed neural networks *J. Comput. Phys.* **394** 136–52
- [21] Daw A, Maruf M and Karpadne A 2021 PID-GAN: a GAN framework based on a physics-informed discriminator for uncertainty quantification with physics *Proc. 27th ACM SIGKDD Conf. on Knowledge Discovery & Data Mining* pp 237–47
- [22] Mises R v 1913 Mechanik der festen körper im plastisch-deformablen zustand *Nachr. Ges. Wiss. Gottingen, Math.-Phys. Kl.* **1913** 582–92
- [23] Erhan D, Courville A, Bengio Y and Vincent P 2010 Why does unsupervised pre-training help deep learning? *Proc. 13th Int. Conf. Artificial Intelligence and Statistics (Proc. Machine Learning Research vol 9)* ed Y W Teh and M Titterton (PMLR) pp 201–8
- [24] Iscen A, Tolias G, Avrithis Y and Chum O 2019 Label propagation for deep semi-supervised learning *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*
- [25] Rabin J, Peyré G, Delon J and Bernot M 2012 Wasserstein barycenter and its application to texture mixing *Scale Space and Variational Methods in Computer Vision* ed A M Bruckstein, B M ter Haar Romeny, A M Bronstein and M M Bronstein (Springer)
- [26] Khrulkov V and Oseledets I 2018 Geometry score: a method for comparing generative adversarial networks *Proc. 35th Int. Conf. on Machine Learning (Proc. Machine Learning Research vol 80)* ed J Dy and A Krause (PMLR) pp 2621–9

- [27] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 Generative adversarial nets *Advances in Neural Information Processing Systems* vol 27, ed Z Ghahramani, M Welling, C Cortes, N Lawrence and K Weinberger (Curran Associates, Inc.)
- [28] Mokhtarishirabad M, Lopez-Crespo P, Moreno B, Lopez-Moreno A and Zanganeh M 2016 Evaluation of crack-tip fields from DIC data: a parametric study *Int. J. Fatigue* **89** 11–19 special Issue: Crack Tip Fields 3
- [29] Radford A, Metz L and Chintala S 2015 Unsupervised representation learning with deep convolutional generative adversarial networks (arXiv:[1511.06434](https://arxiv.org/abs/1511.06434))
- [30] Odena A, Dumoulin V and Olah C 2016 Deconvolution and checkerboard artifacts *Distill* **1**
- [31] Ioffe S and Szegedy C 2015 Batch normalization: accelerating deep network training by reducing internal covariate shift *Int. Conf. on Machine Learning*
- [32] Nair V and Hinton G E 2010 Rectified linear units improve restricted Boltzmann machines *Int. Conf. on Machine Learning (ICML)*
- [33] Maas A L, Hannun A Y and Ng A Y 2013 Rectifier nonlinearities improve neural network acoustic models *Proc. 30th Int. Conf. on Machine Learning*
- [34] Salimans T, Goodfellow I, Zaremba W, Cheung V, Radford A, Chen X and Chen X 2016 Improved techniques for training GANs *Advances in Neural Information Processing Systems* vol 29, ed D Lee, M Sugiyama, U Luxburg, I Guyon and R Garnett (Curran Associates, Inc.)
- [35] Heusel M, Ramsauer H, Unterthiner T, Nessler B and Hochreiter S 2017 GANs trained by a two time-scale update rule converge to a local Nash equilibrium *Advances in Neural Information Processing Systems* vol 30, ed I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan and R Garnett (Curran Associates, Inc.)
- [36] Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V and Rabinovich A 2015 Going deeper with convolutions *2015 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* pp 1–9
- [37] Russakovsky O et al 2015 ImageNet large scale visual recognition challenge *Int. J. Comput. Vis.* **115** 211–52
- [38] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (MIT Press)
- [39] Hatcher A 2000 *Algebraic Topology* (Cambridge University Press)
- [40] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization (arXiv:[1412.6980](https://arxiv.org/abs/1412.6980))
- [41] Arjovsky M and Bottou L 2017 Towards principled methods for training generative adversarial networks (arXiv:[1701.04862](https://arxiv.org/abs/1701.04862))
- [42] Roth K, Lucchi A, Nowozin S and Hofmann T 2017 Stabilizing training of generative adversarial networks through regularization *Proc. 31st Int. Conf. on Neural Information Processing Systems (NIPS 2017)* (Curran Associates Inc.) pp 2015–25
- [43] Arjovsky M, Chintala S and Bottou L 2017 Wasserstein generative adversarial networks *Proc. 34th Int. Conf. on Machine Learning (Proc. of Machine Learning Research* vol 70), ed D Precup and Y W Teh (PMLR) pp 214–23
- [44] Paszke A et al 2019 PyTorch: an imperative style, high-performance deep learning library *Advances in Neural Information Processing Systems* vol 32 (Curran Associates, Inc.) pp 8024–35
- [45] Burt P J and Adelson E H 1987 *The Laplacian Pyramid as a Compact Image Code* (Morgan Kaufmann Publishers Inc.) pp 671–9
- [46] Mescheder L, Nowozin S and Geiger A 2018 Which training methods for GANs do actually converge? *Int. Conf. on Machine Learning (ICML)*
- [47] Mirza M and Osindero S 2014 Conditional generative adversarial nets (arXiv:[1411.1784](https://arxiv.org/abs/1411.1784))