

# GANs for Biological Image Synthesis

Anton Osokin  
INRIA/ENS\*, France  
HSE†, Russia

Anatole Chessel  
École Polytechnique‡,  
France

Rafael E. Carazo Salas  
University of Bristol, UK

Federico Vaggi  
ENS\*, France  
Amazon, USA

## Abstract

*In this paper, we propose a novel application of Generative Adversarial Networks (GAN) to the synthesis of cells imaged by fluorescence microscopy. Compared to natural images, cells tend to have a simpler and more geometric global structure that facilitates image generation. However, the correlation between the spatial pattern of different fluorescent proteins reflects important biological functions, and synthesized images have to capture these relationships to be relevant for biological applications. We adapt GANs to the task at hand and propose new models with casual dependencies between image channels that can generate multi-channel images, which would be impossible to obtain experimentally. We evaluate our approach using two independent techniques and compare it against sensible baselines. Finally, we demonstrate that by interpolating across the latent space we can mimic the known changes in protein localization that occur through time during the cell cycle, allowing us to predict temporal evolution from static images.*

## 1. Introduction

In the life sciences, the last 20 years saw the rise of light fluorescence microscopy as a powerful way to probe biological events in living cells and organisms with unprecedented resolution. The need to analyze quantitatively this deluge of data has given rise to the field of bioimage informatics [29] and is the source of numerous interesting and novel data analysis problems, which current machine learning developments could, in principle, help solve.

Generative models of natural images are among the most long-standing and challenging goals in computer vision. Recently, the community has made significant progress in this task by adopting neural network machinery. Examples of recent models include denoising autoencoders [2], variational autoencoders [20], PixelCNNs [44] and Generative Adversarial Networks (GANs) [14].

\*DI École normale supérieure, CNRS, PSL Research University, Paris

†National Research University Higher School of Economics, Moscow

‡LOB, École Polytechnique, CNRS, INSERM, Université Paris-Saclay

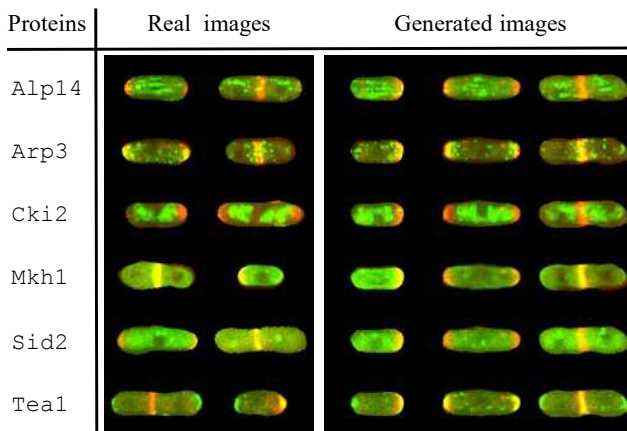


Figure 1. Real (left) and generated (right) images of fission yeast cells with protein Bgs4 depicted in the red channel and 6 other proteins depicted in the green channel. The synthetic images were generated with our star-shaped GAN. The star-shaped model can generate multiple green channels aligned with the same red channel whereas the training images have only one green channel.

GANs [14] are family of successful models, which have recently received widespread attention. Unlike most other generative models, GANs do not rely on training objectives connected to the log likelihood. Instead, GAN training can be seen as a minimax game between two models: the *generator* aims to output images similar to the training set given random noise; while the *discriminator* aims to distinguish the output of the generator from the training set.

Originally, GANs were applied to the MNIST dataset of handwritten digits [21, 14]. The consequent DCGAN model [38] was applied to the CelebA dataset [24] of human faces, the LSUN [52, 38] and ImageNet [7] datasets of natural images. We are not aware of any works applying GANs to biological images.

We work with a recently created bioimage dataset used to extract functional relationships between proteins, called the LIN dataset [9] comprising 170,000 fluorescence microscopy images of cells. In the LIN dataset, each image corresponds to a cell and is composed of signals from two independent fluorescence imaging channels (“red” and “green”), corresponding to the two different proteins tagged with red or green-emitting fluorophores, respectively.

In the LIN dataset, the red channel signal always corresponds to a protein named *Bgs4*, which localizes to the areas of active growth of cells. The green channel signal instead corresponds to any of 41 different “polarity factors”, that is proteins that mark specific areas of the cells’ cortex that help define a cell’s geometry. Polarity factors include proteins like *Alp14*, *Arp3*, *Cki2*, *Mkh1*, *Sid2* or *Tea1* (see Figure 1 for image examples), each of which controls the same biological process “cellular polarity” albeit each in a slightly different way. Each of the green-labeled polarity factors was imaged independently of the others. The biological aim of the LIN study is to investigate how those polarity factors (or proteins) interact with one another.

In this paper, we present a novel application of GANs to generate biological images. Specifically, we want to tackle two concrete limitations of large scale fluorescent imaging screens: we want to use the common information contained in the red channel to learn how to generate a cell with several of the green-labeled proteins together. This would allow us to artificially predict how the localizations of those (independently imaged) proteins might co-exist in cells if they had been imaged together and circumvent the current technical limitations of being able to only image a limited number of signal channels at the same time. Second, taking advantage of the relationship between *Bgs4* and the cell cycle stage, we want to study the dynamical changes in cellular localization that proteins undergo through time as cells grow and divide.

To accomplish this, we make several contributions. We modify the standard DCGAN [38] architecture by substituting the interdependence of the channels with the causal dependence of the green on the red, allowing us to observe multiple modes of green signal for a single red setting. Observing the mode collapse effect of GANs [30, 42] for our separable architecture, we incorporate the recent Wasserstein GAN (WGAN-GP) objective [1, 15]. We propose two approaches to generate multi-channel images: regular WGAN-GP trained on multi-channel images, where extra channels for training are mined by nearest neighbor search in the training set, and a novel star-shaped generator trained directly on the two-channel images. We carefully evaluate our models using two quantitative techniques: the neural network two-sample test (combining ideas from [26] and [15]) and by reconstructing samples in a held out test set with the optimization approach of [30]. For reproducibility, we make the source code and data available online.<sup>1</sup>

This paper is organized as follows. In Section 2, we discuss related works. Section 3 reviews the relevant biological background for our application. In Section 4, we review GANs and present our modeling contributions. We present the experimental evaluation in Section 5 and conclude in Section 6.

<sup>1</sup><https://github.com/aosokin/biogans>

## 2. Related Work

**Generative Adversarial networks (GANs).** Since the seminal paper by Goodfellow *et al.* [14] of 2014 (see also [13] for a detailed review), GANs are becoming an increasingly popular model for learning to generate with the loss functions *learned* jointly with the model itself. Models with adversarial losses have been used in a wide range of applications, such as image generation [8, 38], domain adaptation [12], text-to-image synthesis [39], synthesis of 3D shapes [49] and texture [23], image-to-image translation [18], image super resolution [22] and even generating radiation patterns in particle physics [6]. However, these models suffer from issues such as mode collapse and oscillations during training, making them challenging to use in practice. The community is currently tackling these problems from multiple angles. Extensive effort has been placed on carefully optimizing the architecture of the network [38, 40] and developing best practices to optimize the training procedure.<sup>2</sup> Another active area of research is improving the training objective function [33, 4, 37, 53, 1, 30, 42, 15].

In this paper, we build on the DCGAN architecture [38] combined with the Wasserstein loss [1, 15], where the latter is used to help with the mode collapse issue, appearing especially in our separable setting.

**Conditioning for GANs.** Starting from conditioning on the class labels [31, 8, 34, 11], researchers have extended conditioning to user scribbles [54] and images [46, 18, 55]. While the quality of images generated by [46, 18, 55] is high, their models suffer from conditional mode collapse, i.e., given the first (source) image there is very little or no variety in the second (target image). This effect might be related to the fact that the dataset contained only one target image available for each source image, so the model has only indirect supervision for generating multiple conditioned images. We have applied the pix2pix method of [18] to the LIN dataset and it learned to produce high-quality green images given the red input. However, it was unable to generate multiple realistic green images for one red input.

Given the difficulty in learning robust latent spaces when conditioning on an image, we opted for an alternate approach. We propose a new architecture for the generator, where the red channel and green channels are given independent random noise, and only the red channel is allowed to influence the green channel, see Figure 2 (right).

**Factors of variation.** Chen *et al.* [4] and Mathieu *et al.* [28] used unsupervised methods that encourage disentangling factors of variation in the learned latent spaces, e.g., separating the numerical value of a handwritten digit from its writing style. In contrast to these works, we do not rely on unsupervised training to discover factors of variations, but explicitly embed the separation into the model.

<sup>2</sup><https://github.com/soumith/ganhacks>

**Analysis and synthesis of biological images.** With large scale imaging studies becoming more common in biology, the automated analysis of images is now crucial in many studies to prove the existence of an effect, process large datasets or link with models and simulation [29, 5]. Although the field has only recently embraced deep learning, neural networks are now starting to make a splash, mainly in classical discriminative settings [45].

While, to our knowledge, this work is the first reported use of GANs on samples from fluorescent microscopy, generative models have been widely used in biology [32]. For example, Johnson *et al* [19] learned to generate punctuate patterns in cells (conditional on microtubule localization) showing the potential of those methods in studying the relative sub-cellular positions of several proteins of interest.

Recently, sharing of large biological datasets has greatly improved [25]. Further, EBI has made a large investment to develop the IDR (Image Data Resource) [48], a database built on top of open source tools to facilitate the sharing of terabyte sized datasets with complex metadata.

### 3. Biological Background

#### 3.1. Fluorescent Imaging

Fluorescence microscopy is based on fluorescent compounds, i.e., compounds which can absorb light at given wavelength (the *absorption spectrum*) and re-emit it almost immediately at a slightly different wavelength (the *emission spectrum*). In the case of fluorescent proteins (FPs), of which the Green Fluorescent Protein (GFP) [3, 43] is the first and most widely used one, the fluorescing compound is attached to the protein of interest via genetic engineering. Many FPs of various absorption and emission spectra exist, e.g., Red Fluorescent Protein (RFP) [41]. By genetically tagging different proteins of interest with FPs of different color, one can image them in the same cell at the same time and thus investigate their co-localization. However, the number of proteins that can be tagged and imaged at the same time is limited to 3-4 due to the limited number of FPs with non-overlapping absorption spectra.

Multi-channel fluorescent images are very different from natural images. In natural images, color is determined by the illumination and the properties of a particular material in the scene. In order to generate realistic natural samples, a GAN must capture the relationship between the materials that make up a particular object and its hues. In contrast, in fluorescent images, the intensity of light in a given channel corresponds to the local concentration of the tagged protein, and the correlation between signals in different channels represents important information about the relationship between proteins, but the color does not reflect any intrinsic property about the protein itself.

#### 3.2. Fission Yeast Cells

Fission yeast (*Schizosaccharomyces pombe*) cells are rod shaped unicellular eukaryotes with spherical hemisphere caps. They are born 7  $\mu\text{m}$  long and 4  $\mu\text{m}$  wide, and grow in length to 14  $\mu\text{m}$  while maintaining their width constant. Newly born fission yeast cells start by growing only at the pre-existing end until they reach a critical size, and then switch to bipolar (from the two sides) growth. Bipolar growth continues until cells reach their final length, when they stop growing and start to form a cytokinetic ring in the middle, which is responsible for cleaving the mother cells into two daughters [36]. Interestingly, for most of the cell cycle the length of the cell is a good proxy for its “age”, i.e. the time it has spent growing since its “birth”.

Bgs4, the protein tagged in the red channel, is responsible for cell wall remodeling, and localizes to areas of active growth (see Figure 1 for examples of images). Thus, by observing Bgs4, one can accurately infer growth cycle stage, and predict where cell growth is occurring.

#### 3.3. The LIN Dataset

All experiments in this paper make use of a recent dataset of images of fission yeast cells, which was originally produced to study polarity networks [9]. The LIN dataset consists of around 170,000 of images, with each image being centered on one cell; cell segmentation was performed separately (see [9] for details) and the corresponding outline is also available. Each image is a 3D stack of 2D images where each pixel correspond to a physical size of 100nm; each *z*-plane is distant by 300nm. Every image is composed of two channels, informally called the “red” and the “green”, where light emitted at a precise wavelength is recorded. In this dataset two types of fluorescent-tagged proteins are used: Bgs4 in the red channel, and one of 41 different polarity regulating proteins in the green channel. A full description of all tagged proteins is beyond the scope of this paper: we refer interested readers to [27, 9].

In this paper, we concentrate on a subset of 6 different polarity factors, spanning a large set of different cellular localizations. This gives us 26,909 images of cell, which we, for simplicity, center crop and resize to resolution of 48 $\times$ 80.

### 4. GANs for Image Generation

#### 4.1. Preliminaries

**GAN.** The framework of generative adversarial networks [14, 13] is formulated as a minimax two-player game between two neural networks: *generator* and *discriminator*. The generator constructs images given random noise whereas the discriminator tries to classify if its input image is real (from the training set) or fake (from the generator). The goal of the generator is to trick the discriminator, such

that it cannot easily classify. The discriminator is often referred to as the adversarial loss for training the generator.

More formally, consider a data-generating distribution  $\mathbf{P}_d$  and a training set of images  $\mathbf{x} \in \mathcal{X}$  coming from it. The generator  $G(z; \theta_G)$  is a neural network parameterized by  $\theta_G$  that takes random noise  $z$  from distribution  $\mathbf{P}_z$  as input and produces an image  $\mathbf{x}_{\text{fake}} \in \mathcal{X}$ . The discriminator  $D(\mathbf{x}; \theta_D)$  is a neural network parameterized by  $\theta_D$  that takes either a training image  $\mathbf{x}$  or a generated image  $\mathbf{x}_{\text{fake}}$  and outputs a number in the segment  $[0, 1]$ , where zero is associated with fake images and one – with the real images. As introduced in [14], the key quantity is the negative cross-entropy loss on the discriminator output:

$$L(\theta_D, \theta_G) = \mathbf{E}_{\mathbf{x} \sim \mathbf{P}_{\text{data}}} \log D(\mathbf{x}; \theta_D) + \mathbf{E}_{z \sim \mathbf{P}_z} \log(1 - D(G(z; \theta_G); \theta_D)). \quad (1)$$

The discriminator maximizes (1) w.r.t.  $\theta_D$  and the generator, at the same time, minimizes (1) w.r.t.  $\theta_G$ . In practice, both optimization tasks are attacked simultaneously by alternating between the steps of the two optimizers.

As noted by [14], the objective  $\log(1 - D(G(z; \theta_G); \theta_D))$  often leads to saturated gradients at the initial stages of the training process when the generator is ineffective, i.e., its samples are easy to discriminate from the real data. One practical trick to avoid saturated gradients is to train the generator with maximizing  $\log D(G(z; \theta_G); \theta_D)$  instead.

Goodfellow *et al.* [14] showed that the minimax formulation (1) can be reformulated via minimization of the Jensen-Shannon (JS) divergence<sup>3</sup> between the data-generating distribution  $\mathbf{P}_d$  and the distribution  $\mathbf{P}_G$  induced by  $\mathbf{P}_z$  and  $G$ .

For the architectures of both the generator and the discriminator, we largely reuse a successful version of Radford *et al.* [38] called DCGAN. The generator of DCGAN (see Figure 2, left) is based on up-convolutions [10] interleaved with ReLU non-linearity and batch-normalization [16]. We refer to [38] for additional details.

**Wasserstein GAN.** Recently, Arjovsky *et al.* [1] have demonstrated that in some cases the JS divergence behaves badly and cannot provide any useful direction for training, e.g., when it is discontinuous. To overcome these degeneracies, they consider the earth mover’s distance (equivalent to the 1-st Wasserstein distance) between the distributions

$$W(\mathbf{P}_d, \mathbf{P}_G) = \inf_{\mathbf{P} \in \Pi(\mathbf{P}_d, \mathbf{P}_G)} \mathbf{E}_{(\mathbf{x}, \mathbf{x}') \sim \mathbf{P}} \|\mathbf{x} - \mathbf{x}'\|, \quad (2)$$

where set  $\Pi(\mathbf{P}_d, \mathbf{P}_G)$  is a set of all joint distributions  $\mathbf{P}$  on  $\mathbf{x}$  and  $\mathbf{x}'$  whose marginals are  $\mathbf{P}_d$  and  $\mathbf{P}_G$ , respectively. Intuitively, the distance (2) indicates the cost of the optimal movement of the probability mass from  $\mathbf{P}_d$  to  $\mathbf{P}_G$ . According to [1] by using duality, one can rewrite (2) as

<sup>3</sup>The Jensen-Shannon divergence is a symmetrized version of the Kullback-Leibler divergence between the two distributions, i.e.,  $JS(\mathbf{P}_d, \mathbf{P}_G) = KL(\mathbf{P}_d \| \mathbf{P}_G) + KL(\mathbf{P}_G \| \mathbf{P}_d)$ .

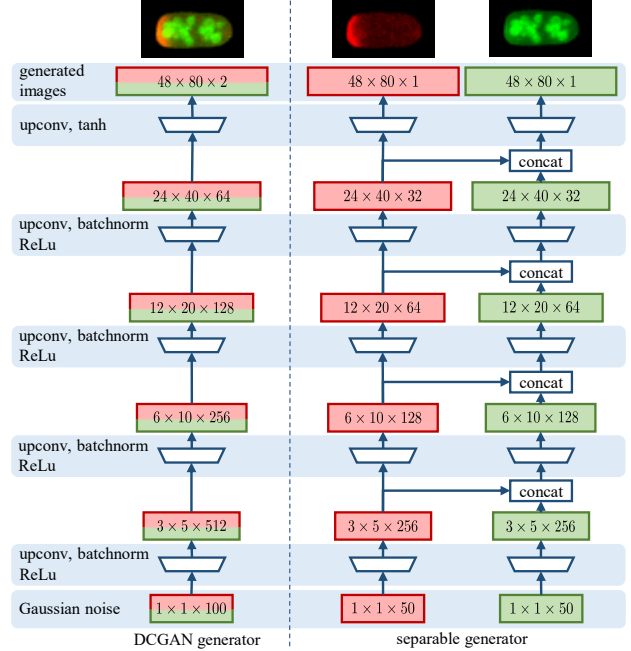


Figure 2. Architectures of the DCGAN generator (left) and our separable generator (right).

$$W(\mathbf{P}_d, \mathbf{P}_G) = \sup_{D \in \mathcal{C}^1} (\mathbf{E}_{\mathbf{x} \sim \mathbf{P}_d} D(\mathbf{x}) - \mathbf{E}_{\mathbf{x}' \sim \mathbf{P}_G} D(\mathbf{x}')), \quad (3)$$

where  $\mathcal{C}^1$  is the set of all 1-Lipschitz functions  $D : \mathcal{X} \rightarrow \mathbb{R}$ .

Optimizing w.r.t. the set  $\mathcal{C}^1$  is complicated. As a practical approximation to the set of all 1-Lipschitz functions, Arjovsky *et al.* [1] suggest to use neural networks  $D(\mathbf{x}; \theta_D)$  with all parameters  $\theta_D$  clipped to a fixed segment. Very recently, Gulrajani *et al.* [15] proposed a surrogate objective to (3), which is based on the  $L_2$ -distance between the norm of the discriminator gradient at specific points and one. In all, we arrive at the minimax game

$$W(\theta_D, \theta_G) = \mathbf{E}_{z \sim \mathbf{P}_z} D(G(z; \theta_G); \theta_D) - \mathbf{E}_{\mathbf{x} \sim \mathbf{P}_{\text{data}}} D(\mathbf{x}; \theta_D) + R(\theta_D), \quad (4)$$

where  $R$  is the regularizer (see [15] for details). The objective (4) is very similar to the original game of GANs (1), but has better convergence properties. In what follows, we refer to the method of [15] as WGAN-GP.

## 4.2. Model Extensions

In this section, we present our modeling contributions. First, we describe our approach to separate the red and green channels of the generator. Second, we discuss a way to train a multi-channel generator using the two-channel data in the LIN dataset. Finally, we propose a new star-shaped architecture that uses the red-green channel separation to obtain multiple channels in the output.

**Channel separation.** The key idea of the channel separation consists in separating the filters of all the up-convolutional layers and the corresponding features into

two halves. The first set of filters is responsible for generating the red channel, while the second half generates the green channel. To make sure the green channel matches the red one, we use one way connections from the red convolutional filters towards the green ones. Figure 2 (right) depicts our modification in comparison to DCGAN (left).

**Multi-channel models.** The LIN dataset [9] contains only two-channel images, the red and one type of the green at a time. Obtaining up to 4 channels simultaneously from a set of 40 proteins (a fixed red and 3 greens) would require the creation of nearly 60,000 yeast strains. Scaling even higher is currently impossible with this imaging technique due to the limited number of FPs with non-overlapping absorption spectra. Because of these constraints, training the generator only on a subset of channels is a task of practical importance. The first approach we present consists in training a multi-channel GAN using an artificial training set of multi-channel images created from the real two-channel images. We proceed as follows: for each two-channel image, we search in every other class for its nearest-neighbors (using  $L_2$ -distance) in the red channel. Then, we create a new sample by combining the original image with the green channels of its nearest neighbors in other classes.

We can then use this dataset to train a multi-output DCGAN. The only difference in the architecture is that the generator outputs  $c+1$  channels, where  $c$  is the number of green channels used in the experiment, and the discriminator takes  $(c+1)$ -channel images as input.

**Star-shaped model.** In our experiments, the multi-channel approach did not perform well, because, even using the nearest neighbors, the extra greens channels were not exactly consistent with the original red signal, emphasizing the importance of correlations between channels.

To overcome this effect, we propose a star-shaped architecture for the generator, consisting of a single red tower (a stack of upconvolutional layers with non-linearities in-between) that feeds into  $c$  green towers (see Figure 2, right). Unlike the multi-channel model described above, the green outputs are independent conditioned on the red. Thus, the model can be trained using the existing two-channel images.

In our experiments, we found it important to use batch normalization [17] in the red tower only once, compared to a more naive way of  $c$  times. The latter leads to interference between several normalizations of the same features and prevents convergence of the training scheme.

After the forward pass, we use  $c$  discriminators attached to different versions of the greens, all paired with the same generated red. For the WGAN-GP version of this model, we apply the original procedure of [15] with the modification that during the discriminator update we simultaneously update all  $c$  discriminators, and the generator receives back the accumulated gradient.

## 5. Experiments

Evaluating generative models is in general non-trivial. In the context of GANs and other likelihood-free approaches, evaluation is even harder, because the models do not provide a way to compute the log-likelihood on the test set, which is the most common evaluation technique. Recently, a number of techniques applicable to evaluating GANs have been proposed [26, 30, 50]. Among those, we chose the following two: the neural-network two-sampled test discussed by [26] combined with the surrogates of the earth mover’s distance [1, 15] and an optimization-based approach of [30] to check if the test samples can be well reconstructed. We modify these techniques to match our needs and check their performance using sensible baselines (Sections 5.1 and 5.2). Finally, in Section 5.3, we show the cell growth cycle generated with our star-shaped model.

### 5.1. Neural-network Two-sample Test

Lopez-Paz and Oquab [26] have recently applied the classifier two-sample test (C2ST) to evaluate the quality of GAN models. A trained generator is evaluated on a held-out test set. This test set is split again into a test-train and test-test subsets. The test-train set is then used to train a fresh discriminator, which tries to distinguish fake images (from the generator) from the real images. Afterwards, the final measure of the quality of the generator is computed as the performance of the new discriminator on the test-test set and the freshly generated images.

When C2ST is applied for images, the discriminator is usually a ConvNet, but even very small ConvNets can discriminate between fake and real images almost perfectly. To obtain a useful measure, Lopez-Paz and Oquab [26] deliberately weaken the ConvNet by fixing some of its parameters to the values obtained by pre-training on ImageNet.

ImageNet-based features are clearly not suitable for LIN cell images, so we weaken the discriminator in another way. We use the negation of the WGAN-GP [15] discriminator objective as a surrogate to the earth mover’s distance. Similar to [26], we train this discriminator on the test-train subset and compute the final estimates on the test-test subset. For all the runs, we repeat the experiment on 10 different random splits of the test set and train the discriminator for 5000 steps with the optimizer used by [15]. For the experiments involving multi-channel generators, we train a separate discriminator for each green channel paired with the red channel.

In our experiments, the training procedure occasionally failed and produced large outliers. To be more robust, we always report a median over 10 random splits together with the median absolute deviation to represent the variance. In the Suppl. Mat. [35], we additionally quantitatively and qualitatively compare the WGAN-GP [15], WGAN [1] and cross-entropy discriminators used in C2ST.

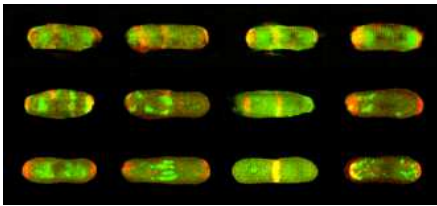
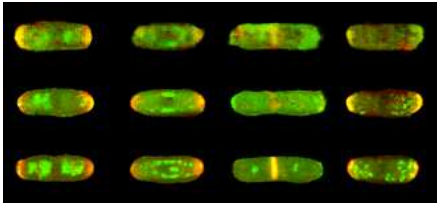
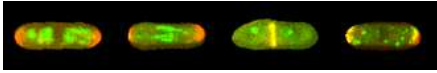
	Steps	C2ST	Samples
sep. GAN	1k	$11.0 \pm 0.1$	
	5k	$6.7 \pm 0.1$	
	50k	$3.2 \pm 0.1$	
sep. WGAN-GP	1k	$6.0 \pm 0.1$	
	5k	$2.2 \pm 0.1$	
	50k	$1.6 \pm 0.1$	
Real	-	$-0.7 \pm 0.6$	

Figure 3. Scores of the classifier two-sample test (C2ST) between the generators and the hold-out test sets of images. We report the scores of separable GAN and WGAN-GP at different stages of training. For each line, we show the samples from the corresponding models to demonstrate that the lower C2ST scores correspond to better-looking (sharper, less artifacts, etc.) images. Best viewed in color and on a screen. An extended version of this figure is given in the Suppl. Mat. [35].

**Sanity checks of the two-sample test.** We evaluate C2ST in two baseline settings. First, we compare the separable GAN [38] and the WGAN-GP [15] models (based on the same DCGAN architecture, trained on the same set of images of 6 proteins) at different stages of the training process. For each of these models, we also show qualitative difference between the generated images. Figure 3 shows that along the training process, quality of both GAN and WGAN-GP improves, i.e., generated images become sharper and contain less artifacts, consistent with the C2ST score. To better visualize the difference between the trained GAN and WGAN-GP models, in Figure 4, we show multiple samples of the green channel corresponding to the same red channel. We see that the C2ST evaluation captures several aspects of the visual quality (such as sharpness, correct shape, absence of artifacts, diversity of samples) and provides a meaningful score.

From Figures 3 and 4, we also conclude that the quality of GAN samples is worse than the quality of WGAN-GP according to visual inspection. C2ST (based on WGAN-GP) confirms this observation, which is not surprising given that WGAN-GP was trained using the same methodology. Surprisingly, when evaluated with the cross-entropy C2ST, WGAN-GP also performs better than GAN (see Suppl. Mat. [35] for details).

As the second baseline evaluation, we use C2ST to compare real images of different classes. Table 1 shows that when evaluated w.r.t. the test set of the same class the esti-

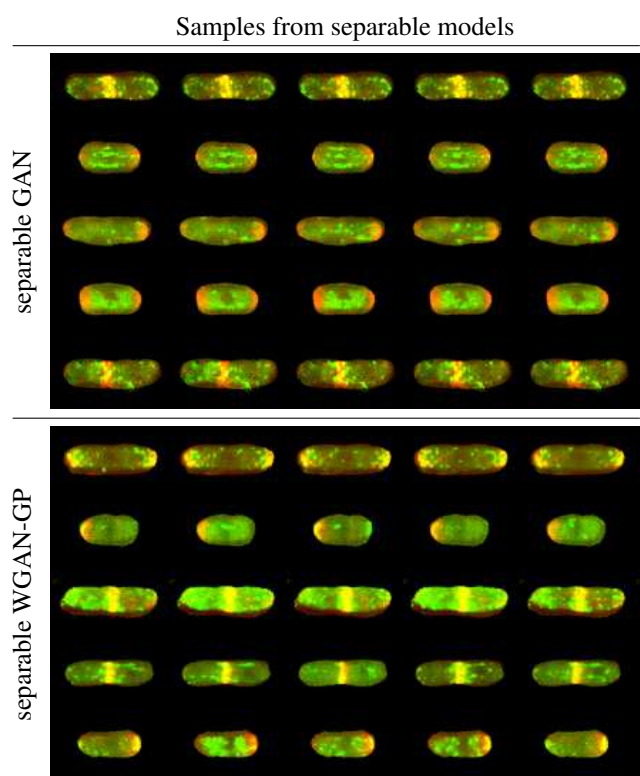


Figure 4. Samples generated by separable GAN (top) and WGAN-GP (bottom) models trained on the 6 selected proteins shown in Figure 1. Each row has samples with identical red channel, but different green ones. We observe that WGAN-GP provides much larger variability of the green channel conditioned on the red. In particular, in the three bottom rows, even the type of the protein changes, which we have never observed for the samples of GAN (this effect should be present, because the model is trained without any distinction between the classes, but is surprisingly rare). This difference is captured by the C2ST evaluation: the GAN model has the score of  $3.2 \pm 0.1$  compared to  $1.6 \pm 0.1$  of WGAN-GP.

mates are significantly smaller (but with non-zero variance) compared to when evaluated w.r.t. different classes. Note that the C2ST score is not a metric. In particular, Table 1 is not symmetric reflecting biases between the train/test splits. Specifically to WGAN-GP, the score can also be negative, because the quadratic regularization term is the dominant part of the objective (4) when the two image sources are very similar.

As an additional test, we include two extra proteins *Fim1* and *Tea4* that are known to have similar localization to *Arp3* and *Tea1*, respectively. We observe that C2ST reflects this similarity by giving the pairs of similar proteins a much smaller score compared to most of other pairs (but still significantly higher than comparing a protein to itself).

**Results.** Table 2 shows the results of C2ST applied to several models with multiple output channels (see Section 4.2): the multi-channel model and its separable version, the star-shaped model and the two baselines, which do not align

test	Alp14	Arp3	Cki2	Mkh1	Sid2	Tea1	Fim1	Tea4
Alp14	$0.1 \pm 0.2$	$12.5 \pm 0.3$	$8.1 \pm 0.3$	$12.5 \pm 0.5$	$9.5 \pm 0.2$	$10.9 \pm 0.3$	$15.6 \pm 0.3$	$11.4 \pm 0.3$
Arp3	$14.4 \pm 0.2$	$0.8 \pm 0.4$	$16.2 \pm 0.2$	$11.5 \pm 0.4$	$20.5 \pm 0.3$	$13.2 \pm 0.2$	<b><math>3.7 \pm 0.2</math></b>	$18.3 \pm 0.3$
Cki2	$8.6 \pm 0.2$	$15.9 \pm 0.3$	$-0.2 \pm 0.3$	$13.7 \pm 0.4$	$12.0 \pm 0.3$	$15.8 \pm 0.3$	$18.5 \pm 0.4$	$16.0 \pm 0.5$
Mkh1	$12.3 \pm 0.4$	$12.2 \pm 0.6$	$13.6 \pm 0.3$	$-0.2 \pm 0.4$	$12.4 \pm 0.6$	$13.3 \pm 0.6$	$15.1 \pm 0.5$	$14.9 \pm 0.8$
Sid2	$9.0 \pm 0.3$	$19.5 \pm 0.4$	$11.8 \pm 0.5$	$13.4 \pm 0.9$	$-0.6 \pm 0.3$	$12.6 \pm 0.3$	$23.9 \pm 0.4$	$7.7 \pm 0.6$
Tea1	$11.3 \pm 0.3$	$11.5 \pm 0.5$	$15.9 \pm 0.3$	$14.4 \pm 0.6$	$13.1 \pm 0.1$	$-0.1 \pm 0.4$	$14.5 \pm 0.5$	<b><math>6.9 \pm 0.5</math></b>
Fim1	$16.3 \pm 0.2$	<b><math>2.8 \pm 0.3</math></b>	$18.4 \pm 0.2$	$14.5 \pm 0.3$	$23.4 \pm 0.3$	$15.1 \pm 0.2$	$-0.2 \pm 0.3$	$20.8 \pm 0.5$
Tea4	$9.7 \pm 0.6$	$15.8 \pm 0.7$	$14.0 \pm 0.9$	$13.9 \pm 0.9$	$6.2 \pm 0.4$	<b><math>5.9 \pm 0.3</math></b>	$19.5 \pm 0.7$	$-0.5 \pm 0.7$

Table 1. Results of C2ST with WGAN-GP when comparing real images of different proteins. For each run, the training images of one class are evaluated w.r.t. the test images of another class. The reported values are comparable with Table 2, but not with Figure 3.

	real images	one-class non-separable	one-class separable	multi-channel non-separable	multi-channel separable	star-shaped
separable red/green class conditioned	-	$\times$	$\checkmark$	$\times$	$\checkmark$	$\checkmark$
	-	$\times$	$\times$	$\checkmark$	$\checkmark$	$\checkmark$
Alp14	$0.1 \pm 0.2$	<b><math>0.6 \pm 0.3</math></b>	$1.2 \pm 0.2$	$3.2 \pm 0.4$	$2.3 \pm 0.5$	<b><math>0.6 \pm 0.3</math></b>
Arp3	$0.8 \pm 0.4$	<b><math>1.2 \pm 0.3</math></b>	$2.4 \pm 0.4$	$3.2 \pm 0.4$	$4.2 \pm 0.4$	$2.1 \pm 0.5$
Cki2	$-0.2 \pm 0.3$	<b><math>0.3 \pm 0.5</math></b>	$1.0 \pm 0.3$	$2.5 \pm 0.3$	$3.6 \pm 0.5$	$1.2 \pm 0.3$
Mkh1	$-0.2 \pm 0.4$	$0.8 \pm 0.6$	<b><math>0.5 \pm 0.4</math></b>	$4.6 \pm 0.5$	$6.6 \pm 0.5$	$2.4 \pm 0.6$
Sid2	$-0.6 \pm 0.3$	<b><math>0.8 \pm 0.4</math></b>	$1.0 \pm 0.5$	$4.5 \pm 0.5$	$3.2 \pm 0.6$	$1.1 \pm 0.6$
Tea1	$-0.1 \pm 0.4$	<b><math>0.8 \pm 0.5</math></b>	<b><math>0.8 \pm 0.5</math></b>	$4.4 \pm 0.3$	$2.8 \pm 0.5$	$1.1 \pm 0.4$
6 proteins	$-0.1 \pm 0.2$	<b><math>0.8 \pm 0.2</math></b>	$1.1 \pm 0.2$	$3.7 \pm 0.1$	$3.8 \pm 0.2$	$1.4 \pm 0.1$

Table 2. Results of C2ST with the WGAN-GP objective comparing several multi-channel models w.r.t. the real images. All the models were trained with WGAN-GP. The values in this table are directly comparable to the ones in Table 1.

green channels of different classes with the same red channel: one-class generators trained individually for each class and their separable versions. All the models were trained with WGAN-GP with the same ratio of the width of the generator tower to the number of output channels.

We observe that the individual one-class WGAN-GP models lead to higher quality compared to all the models outputting synchronized channels for all the classes. Among the models that synchronize channels, the star-shaped model performs best, but for some proteins there is a significant drop in quality w.r.t. the one-class models.

## 5.2. Optimization to Reconstruct the Test Set

One of the common failures of GANs is the loss of modes from the distribution, usually referred to as mode collapse. There is evidence [37] that image quality can be inversely correlated with mode coverage. To test for the mode collapse, we perform an experiment proposed in [30], where for a fixed trained generator  $G$  we examine how well it can reconstruct images from a held out test set. For each image in the test set, we minimize the  $L_2$ -distance (normalized by the number of pixels) between the generated and test images w.r.t. the noise vector  $z$ . We call this task regular reconstruction. We use 50 iterations of L-BFGS and run it 5 times to select the best reconstruction. We also performed an additional task, separable reconstruction, which examines the ability of separable networks to reproduce modes of the green channel conditioned on the red. In this task, we

use a two-step procedure: first, we minimize the  $L_2$ -error between the red channels holding the green noise fixed, and then we minimize the  $L_2$ -error in the green channel while keeping the red noise fixed at its optimized value. To complete the study, we also report the negative log likelihood (NLL) w.r.t. the prior  $P_z$  of the noise vectors  $z$  obtained with a reconstruction procedure. As a baseline for the reconstruction error, we show the nearest neighbor cell (in both red and green channels) from the training set and the average  $L_2$ -distance to the nearest neighbors. As a baseline for NLL, we show the averaged NLL for the random point generated from  $P_z$ .

We apply the reconstruction procedure to evaluate four models: separable one-class and star-shaped models trained with both GAN and WGAN-GP algorithms. Figure 5 and Table 3 present qualitative and quantitative results, respectively. For all the measurements, we report the median values and the median absolute deviation. In Figure 6, we plot reconstruction errors vs. NLL values for the Mkh1, which was the hardest protein in the separable reconstruction task.

Analyzing the results, we observe that separable reconstruction is a harder task than the single step procedure. Second, WGAN-GP models can reconstruct better, probably because they suffer less from the mode collapse. And finally, the star-shaped models do not degrade the performance in terms of reconstruction, except for some hard proteins (see more details in the Suppl. Mat. [35]).

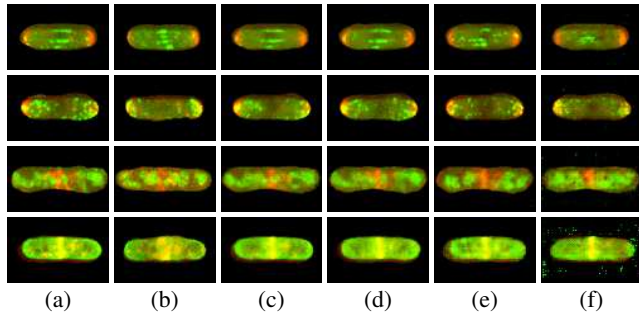


Figure 5. Examples of cell reconstructions. (a) – a test image; (b) – the  $L_2$  nearest neighbor; (c) – regular reconstruction by one-class separable WGAN-GP; (d) – regular reconstruction by star-shaped WGAN-GP; (e) – separable reconstruction by star-shaped WGAN-GP; (f) – separable reconstruction by star-shaped GAN. An extended version of this figure is given in the Suppl. Mat. [35].

		$L_2$ -error	NLL
regular	Nearest neighbors	$0.079 \pm 0.009$	-
	Gaussian noise	-	$142 \pm 5$
	GAN-sep	$0.053 \pm 0.007$	$166 \pm 17$
	WGAN-GP-sep	$0.043 \pm 0.006$	$149 \pm 8$
separable	GAN-star	$0.061 \pm 0.008$	<b><math>139 \pm 12</math></b>
	WGAN-GP-star	<b><math>0.041 \pm 0.005</math></b>	$150 \pm 8$
	GAN-sep	$0.069 \pm 0.011$	$158 \pm 13$
	WGAN-GP-sep	$0.062 \pm 0.009$	$143 \pm 6$
	GAN-star	$0.074 \pm 0.011$	<b><math>142 \pm 7</math></b>
	WGAN-GP-star	<b><math>0.058 \pm 0.010</math></b>	$143 \pm 7$

Table 3. Reconstruction experiment. For the four trained models (GAN/WGAN-GP and separable one-class/star-shaped), we report  $L_2$ -errors of the reconstructions and the negative log likelihoods (NLL) of the latent vectors found by the reconstruction.

### 5.3. Progression Through the Cell Cycle

As described in Section 3.2, the localization of Bgs4 can be used to accurately pinpoint the cell cycle stage. However, not nearly as much as is known about how the localization of the other proteins changes within the cell cycle [27].

Using our separable GAN architecture, we can interpolate between points in the latent space [47] to move across the different stages of growth and division. Due to the architecture of our network, the output of the green channel will always remain consistent with the red output. We show an example of the reconstructed cell cycle in Figure 7 and several animated examples in the Suppl. Mat. [35]. As a validation of our approach, Arp3 is seen gradually moving a dot like pattern at the tips of the cell towards the middle of the cell during mitosis, as has been previously described in the literature [51].

It’s important to highlight that the LIN dataset lacks true multi-channel (3+) images, and as such, we are unable to compare how our generated multi-channel images compare to real fluorescent images. We hope that as more datasets

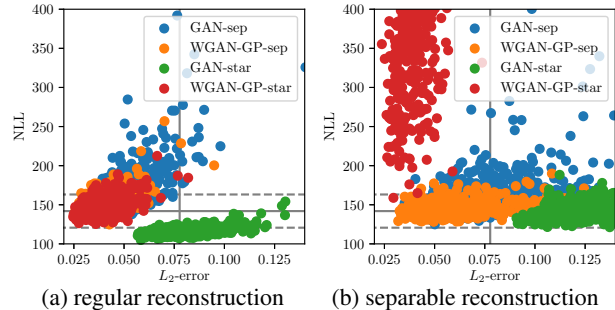


Figure 6. Reconstruction errors against negative log likelihood (NLL) of the latent vectors found by reconstruction. We show all the cells corresponding to protein Mkh1, which appears to be the hardest for the star-shaped models. The vertical gray line shows the median  $L_2$ -error of the nearest neighbor. Horizontal gray lines show mean NLL ( $\pm 3$  std) of the noise sampled from the Gaussian prior. In the separable (red-first) setting, the star-shaped model trained with GAN provides very bad reconstructions, whereas the same model trained with WGAN-GP results in high NLL values. An extended version of this figure is given in the Suppl. Mat. [35].

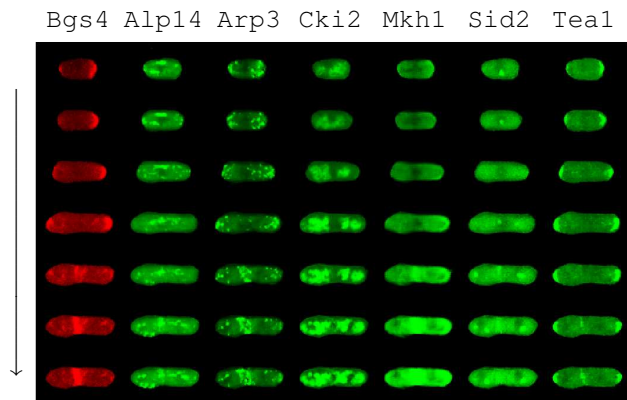


Figure 7. Cell cycle of a star-shaped WGAN-GP model.

in biology become open, we will have a better baseline to compare our model too.

## 6. Conclusion

Although generative modeling has seen an explosion in popularity in the last couple of years, so far it has mostly been applied to the synthesis of real world images. Our results in this paper suggest that modern generative models can be fruitfully applied to images obtained by fluorescent microscopy. By leveraging correlation between different image channels, we were able to simulate the localization of multiple proteins throughout the cell cycle. This could enable in the future the exploration of uninvestigated, inaccessible or unaffordable biological/biomedical experiments, to catalyze new discoveries and potentially enable new diagnostic and prognostic bioimaging applications.

**Acknowledgements.** A. Osokin was supported by the ERC grant Activia (no. 307574). F. Vaggi was supported by a grant from the chair Havas-Dauphine.

## References

- [1] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017. 2, 4, 5
- [2] Y. Bengio, L. Yao, G. Alain, and P. Vincent. Generalized denoising auto-encoders as generative models. In *NIPS*, 2013. 1
- [3] M. Chalfie, Y. Tu, G. Euskirchen, W. W. Ward, and D. C. Prasher. Green fluorescent protein as a marker for gene expression. *Science*, pages 802–805, 1994. 3
- [4] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*, 2016. 2
- [5] A. Chessel. An overview of data science uses in bioimage informatics. *Methods*, 2017. 3
- [6] L. de Oliveira, M. Paganini, and B. Nachman. Learning particle physics by example: Location-aware generative adversarial networks for physics synthesis. *arXiv:1701.05927v1*, 2017. 2
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 1
- [8] E. L. Denton, S. Chintala, A. Szlam, and R. Fergus. Deep generative image models using a Laplacian pyramid of adversarial networks. In *NIPS*, 2015. 2
- [9] J. Dodgson, A. Chessel, F. Vaggi, M. Giordan, M. Yamamoto, K. Arai, M. Madrid, M. Geymonat, J. F. Abenza, J. Cansado, M. Sato, A. Csikasz-Nagy, and R. E. Carazo Salas. Reconstructing regulatory pathways by systematically mapping protein localization interdependency networks. *bioRxiv:11674*, 2017. 1, 3, 5
- [10] A. Dosovitskiy, J. T. Springenberg, M. Tatarchenko, and T. Brox. Learning to generate chairs, tables and cars with convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 39:692–705, 2017. 4
- [11] V. Dumoulin, I. Belghazi, B. Poole, O. Mastropietro, A. Lamb, M. Arjovsky, and A. Courville. Adversarially learned inference. In *ICLR*, *arXiv:1606.00704v3*, 2017. 2
- [12] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016. 2
- [13] I. Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *arXiv:1701.00160v3*, 2017. 2, 3
- [14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 1, 2, 3, 4
- [15] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. Improved training of Wasserstein GANs. *arXiv:1704.00028v2*, 2017. 2, 4, 5, 6
- [16] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 4
- [17] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 5
- [18] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 2
- [19] G. R. Johnson, J. Li, A. Shariff, G. K. Rohde, and R. F. Murphy. Automated Learning of Subcellular Variation among Punctate Protein Patterns and a Generative Model of Their Relation to Microtubules. *PLoS computational biology*, 11(12):e1004614, 2015. 3
- [20] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *ICLR*, *arXiv:1312.6114v10*, 2014. 1
- [21] Y. LeCun, C. Cortes, and C. J. C. Burges. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998. 1
- [22] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017. 2
- [23] C. Li and M. Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. In *ECCV*, 2016. 2
- [24] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 1
- [25] V. Ljosa, K. L. Sokolnicki, and A. E. Carpenter. Annotated high-throughput microscopy image sets for validation. *Nature Methods*, 9(7):637, 2012. 3
- [26] D. Lopez-Paz and M. Oquab. Revisiting classifier two-sample tests. In *ICLR*, *arXiv:1610.06545v3*, 2017. 2, 5
- [27] S. G. Martin and R. A. Arkowitz. Cell polarization in budding and fission yeasts. *FEMS microbiology reviews*, 38(2):228–253, 2014. 3, 8
- [28] M. Mathieu, J. Zhao, P. Sprechmann, A. Ramesh, and Y. LeCun. Disentangling factors of variation in deep representation using adversarial training. In *NIPS*, 2016. 2
- [29] E. Meijering, A. E. Carpenter, H. Peng, F. A. Hamprecht, and J.-C. Olivo-Marin. Imagining the future of bioimage analysis. *Nature Biotechnology*, 34(12):1250–1255, 2016. 1, 3
- [30] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein. Unrolled generative adversarial networks. In *ICLR*, *arXiv:1611.02163v3*, 2017. 2, 5, 7
- [31] M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv:1411.1784v1*, 2014. 2
- [32] R. F. Murphy. Building cell models and simulations from microscope images. *Methods*, 96:33–39, 2016. 3
- [33] S. Nowozin, B. Cseke, and R. Tomiokao. f-GAN: Training generative neural samplers using variational divergence minimization. In *NIPS*, 2016. 2
- [34] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier GANs. In *ICML*, 2017. 2
- [35] A. Osokin, A. Chessel, R. E. Carazo Salas, and F. Vaggi. GANs for biological image synthesis. Code and supplementary materials. <https://github.com/aosokin/biogans>, 2017. 5, 6, 7, 8
- [36] T. D. Pollard and J.-Q. Wu. Understanding cytokinesis: lessons from fission yeast. *Nature reviews Molecular cell biology*, 11(2):149–155, 2010. 3
- [37] B. Poole, A. A. Alemi, J. Sohl-Dickstein, and A. Angelovam. Improved generator objectives for GANs. In *NIPS Workshop on Adversarial Training*, *arXiv:1612.02780v1*, 2016. 2, 7

- [38] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, *arXiv:1511.06434v2*, 2016. 1, 2, 4, 6
- [39] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text to image synthesis. In *ICML*, 2016. 2
- [40] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. In *NIPS*, 2016. 2
- [41] N. C. Shaner, R. E. Campbell, P. A. Steinbach, B. N. Giepmans, A. E. Palmer, and R. Y. Tsien. Improved monomeric red, orange and yellow fluorescent proteins derived from *Discosoma sp.* red fluorescent protein. *Nature biotechnology*, 22(12):1567, 2004. 3
- [42] I. Tolstikhin, S. Gelly, O. Bousquet, C.-J. Simon-Gabriel, and B. Schölkopf. AdaGAN: Boosting generative models. *arXiv:1701.02386v1*, 2017. 2
- [43] R. Y. Tsien. The green fluorescent protein. *Annual review of biochemistry*, 67:509–544, 1998. 3
- [44] A. van den Oord, N. Kalchbrenner, L. Espeholt, K. Kavukcuoglu, O. Vinyals, and A. Graves. Conditional image generation with PixelCNN decoders. In *NIPS*, 2016. 1
- [45] D. A. Van Valen, T. Kudo, K. M. Lane, D. N. Macklin, N. T. Quach, M. M. DeFelice, I. Maayan, Y. Tanouchi, E. A. Ashley, and M. W. Covert. Deep Learning Automates the Quantitative Analysis of Individual Cells in Live-Cell Imaging Experiments. *PLOS Computational Biology*, 12(11):e1005177, 2016. 3
- [46] X. Wang and A. Gupta. Generative image modeling using style and structure adversarial networks. In *ECCV*, 2016. 2
- [47] T. White. Sampling generative networks. *arXiv:1609.04468v3*, 2016. 8
- [48] E. Williams, J. Moore, S. W. Li, G. Rustici, A. Tarkowska, A. Chessel, S. Leo, B. Antal, R. K. Ferguson, U. Sarkans, A. Brazma, R. E. Carazo Salas, and J. Swedlow. The Image Data Resource: A scalable platform for biological image data access, integration, and dissemination. *Nature Methods*, 14:775–781, 2017. 3
- [49] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *NIPS*, 2016. 2
- [50] Y. Wu, Y. Burda, R. Salakhutdinov, and R. Grosse. On the quantitative analysis of decoder-based generative models. In *ICLR*, *arXiv:1611.04273v1*, 2017. 5
- [51] H. Yan and M. K. Balasubramanian. Meiotic actin rings are essential for proper sporulation in fission yeast. *Journal of Cell Science*, 125(6):1429–1439, 2012. 8
- [52] 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:1506.03365v3*, 2016. 1
- [53] J. Zhao, M. Mathieu, and Y. LeCun. Energy-based generative adversarial networks. In *ICLR*, *arXiv:1609.03126v4*, 2017. 2
- [54] J.-Y. Zhu, P. Krähenbühl, E. Shechtman, and A. A. Efros. Generative visual manipulation on the natural image manifold. In *ECCV*, 2016. 2
- [55] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 2