

Focus on Content not Noise: Improving Image Generation for Nuclei Segmentation by Suppressing Steganography in CycleGAN

Jonas Utz¹ Tobias Weise^{1,2} Maja Schlereth¹ Fabian Wagner²
Mareike Thies² Mingxuan Gu² Stefan Uderhardt^{3,4,5} Katharina Breininger¹

¹Department Artificial Intelligence in Biomedical Engineering, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

² Pattern Recognition Lab, FAU Erlangen-Nürnberg, Germany

³ Department of Medicine 3 - Rheumatology and Immunology, FAU Erlangen-Nürnberg, Germany

⁴ Deutsches Zentrum für Immuntherapie, FAU Erlangen-Nürnberg, Germany

⁵ Exploratory Research Unit, Optical Imaging Centre Erlangen, FAU Erlangen-Nürnberg, Germany

jonas.utz@fau.de

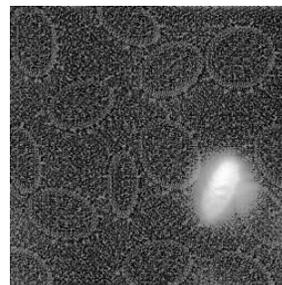
Abstract

Annotating nuclei in microscopy images for the training of neural networks is a laborious task that requires expert knowledge and suffers from inter- and intra-rater variability, especially in fluorescence microscopy. Generative networks such as CycleGAN can inverse the process and generate synthetic microscopy images for a given mask, thereby building a synthetic dataset. However, past works report content inconsistencies between the mask and generated image, partially due to CycleGAN minimizing its loss by hiding shortcut information for the image reconstruction in high frequencies rather than encoding the desired image content and learning the target task. In this work, we propose to remove the hidden shortcut information, called steganography, from generated images by employing a low pass filtering based on the discrete cosine transform (DCT). We show that this increases coherence between generated images and cycled masks and evaluate synthetic datasets on a downstream nuclei segmentation task. Here we achieve an improvement of 5.4 percentage points in the F1-score compared to a vanilla CycleGAN. Integrating advanced regularization techniques into the CycleGAN architecture may help mitigate steganography-related issues and produce more accurate synthetic datasets for nuclei segmentation.

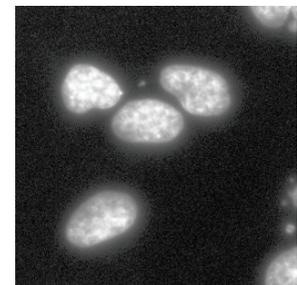
1. Introduction

Detecting individual nuclei and cellular structures in microscopy images is a common task in biomedical image analysis and provides insights into cellular processes. Deep neural networks have been successfully used to au-

tomate nuclei detection but are typically trained in a supervised manner with a paired dataset of nuclei images and manual labels [8, 10]. However, acquiring manual labels is time-consuming and prone to errors and inconsistencies, especially in fluorescence microscopy (FM) images, e.g., due to clustered nuclei, bleed-through, and other artifacts. Unsupervised methods present an alternative, with two approaches being commonly used: a two-stage approach where a generative model is used to generate a synthetic dataset for training a supervised segmentation network and a one-stage approach where a generative model is directly used to perform nuclei segmentation. In both cases, the generative model is trained with an unpaired dataset of real microscopy images and synthetic masks which can be obtained from a suitable synthesis model such as ellipses [1]. The model is then trained to perform unpaired image-to-image translation between masks and nuclei images. A



(a) Generated nuclei



(b) Real nuclei

Figure 1: 1a image with steganography amplified by adaptive histogram equalization. The generated image contains information in seemingly empty regions. 1b Processed real nuclei image for reference.

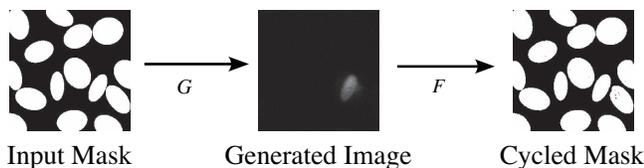


Figure 2: Severe example of CycleGAN steganography: Despite only one nucleus being visible in the generated image, generator F can reconstruct the input mask almost perfectly. Generator G encodes information about the input mask in the generated image in the form of a high-frequency low-amplitude signal. This signal is used by F to reconstruct the original mask.

commonly chosen approach is the CycleGAN architecture, introduced by Zhu *et al.* [13]. A CycleGAN employs two generators, each dedicated to translate between two image domains and two discriminators for the respective domains. A cycle-consistency loss for the generators enforces consistency between an original image and its cycled version to ensure the learned mapping is a bijection. Essentially, the *cycle loss* ensures that the generators retain the image content and only transfer the domain (*e.g.*, from mask to nuclei). However, this constraint does not necessarily lead to a content similarity between the input and the generated image in the other domain. As a shortcut to minimize the cycle-consistency loss, generators can embed a compressed version of the input in the generated image as a high-frequency low-amplitude signal to create a perfect reconstruction of the input image [4, 7]. This emerging property has been denoted as *CycleGAN steganography* since the generator actively hides information from the discriminator. This can cause a mismatch between the source image and the generated image leading to label noise in the context of synthetic dataset generation. The prominent effect observed in nuclei images could be attributed to the minimal influence of missing instances on the discriminator’s output. Since the discriminator’s role is to classify whether an image appears real, the absence of certain nuclei instances does not considerably impact the realism of the microscopy image. An impression of the CycleGAN steganography effect can be obtained from Figure 2. We visualize the steganography signal in Figure 1, by performing adaptive histogram equalization, which is an image processing technique that enhances contrast by redistributing pixel intensities based on local regions. For reference, we apply the adaptive histogram equalization not only to the generated image but also a real nuclei image. The unprocessed version of Figure 1a is also used in Figure 2. The CycleGAN steganography property was first reported by Chu *et al.* in [4]. Since the hidden signal is of low amplitude with a high frequency, it has characteristics similar to noise. Motivated by this observation, Porav *et al.* [7] pro-

pose using DnCNN, an end-to-end differentiable image denoising network, between the two generators to remove hidden information. The denoising network is only applied for the cycle-consistency loss and is frozen during CycleGAN training. They evaluate their method on a day-to-night style transfer task with semantic segmentation. Löhdefink and Fingscheidt [5] try to corrupt the hidden signal and break undesired information exchange between the generators by adding Gaussian noise or discretization noise to the generated images. They evaluate their approach on a supervised semantic segmentation of driving scenes.

In this paper, we generate a fully synthetic dataset of FM images of nuclei with a CycleGAN using transformed ellipses as input masks. To mitigate the before-mentioned CycleGAN steganography problem, we propose to use the discrete cosine transform (DCT) to perform low-pass filtering of the generated images to remove the hidden signal adaptively. While we still observe a discrepancy in content between the input mask and generated nuclei images using our approach, the cycled mask matches the content of the generated image. This allows us to use the generated images with their cycled masks as training data for the supervised training of segmentation networks. We evaluate our approach by testing only with synthetic data trained networks on real data. We observe an increased average precision when DCT filtering is used due to the improved alignment between the mask and generated image. To summarize, the contributions of our work are:

- We compare different CycleGAN steganography removal techniques on FM images.
- We propose an efficient way to remove steganography from CycleGAN images by introducing an additional filter step, based on the DCT.
- We analyze the impact of CycleGAN steganography on downstream segmentation performance on two public FM nuclei datasets and show improved performance when using our proposed method.

2. Methods

2.1. CycleGAN

Our method is based on CycleGAN[13] which can learn a translation between two domains, such as the domain of masks X and nuclei images Y , with individual samples from the domains denoted as x and y respectively. This translation is performed by two generator networks $G : X \rightarrow Y$ and $F : Y \rightarrow X$. Both generators are trained with an adversarial loss using discriminator networks D_X and D_Y for each domain. For generator G , which receives a mask image x and generates a fake nuclei image \hat{y} , the

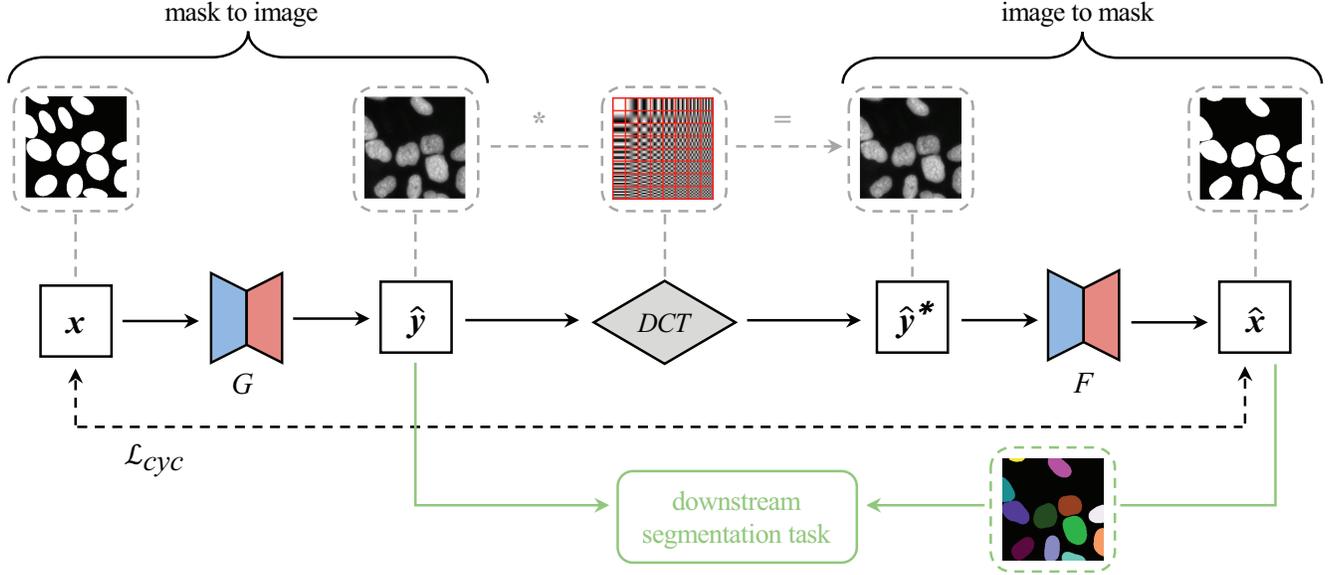


Figure 3: Extended CycleGAN architecture featuring a DCT-based low-pass filter, effectively removing undesired steganography from generated images. The downstream segmentation network evaluates the model’s performance and is trained on generated images and cycled masks.

adversarial loss \mathcal{L}_{adv} is defined as:

$$\mathcal{L}_{adv} = \mathbb{E}_{\mathbf{y} \sim p_{data}(\mathbf{y})} [\log D_Y(\mathbf{y})] + \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log (1 - D_Y(G(\mathbf{x})))] \quad (1)$$

The adversarial loss for generator F is computed analogously. An additional *cycle consistency loss* is used to ensure that the mapping is bijective by enforcing $F(G(\mathbf{x})) \approx \mathbf{x}$ and $G(F(\mathbf{y})) \approx \mathbf{y}$. In other words, a sample should be close to its original after being consecutively processed by both generators. The cycle-consistency loss for generator G (F again analogously) is formulated as:

$$\mathcal{L}_{cyc} = \mathbb{E}_{\mathbf{y} \sim p_{data}(\mathbf{y})} [\|G(F(\mathbf{y})) - \mathbf{y}\|_1]. \quad (2)$$

2.2. DCT filtering

The DCT algorithm is best known in the context of the JPEG image compression algorithm [11], which removes high frequencies from images for compression. We propose eliminating the high-frequency pixel-level patterns [4] used by CycleGANs to hide information by setting coefficients associated with higher frequencies to zero. The DCT is a bijective mapping between an input I and the frequency domain. In this work, I refers to the generated image by CycleGAN, *i.e.*, $\hat{\mathbf{y}}$ or $\hat{\mathbf{x}}$. In the following equations, I_{xy} represents the pixel intensity at position (x, y) in the image. The coefficients C of the DCT are computed on the image

I of $N \times M$ pixels as follows:

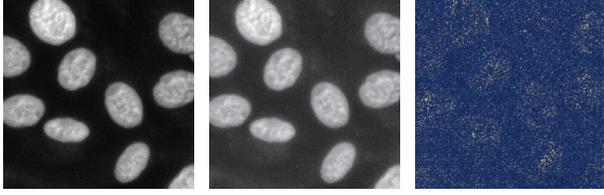
$$\text{DCT}(i, j) = \alpha(i)\alpha(j) \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} I_{xy} h v, \quad (3)$$

$$\text{where } h = \cos \frac{\pi(2x+1)i}{2N}, v = \cos \frac{\pi(2y+1)j}{2M}.$$

The functions h and v weight the contributions of different frequency components in the image I_{xy} to calculate the DCT coefficient at position (i, j) . The scaling factor α can be used to make the transform orthonormal:

$$\alpha(k) = \begin{cases} \frac{1}{\sqrt{2}}, & \text{if } k = 0 \\ 1, & \text{otherwise} \end{cases}. \quad (4)$$

When computed on $2D$ images of size $N \times M$, $N \times M$ coefficients $C^{M \times N}$ are obtained. In order to eliminate specific frequency bands from the image, we applied a technique of setting corresponding DCT coefficients to zero in the frequency domain. This was achieved by element-wise multiplication of the DCT coefficients matrix with a mask matrix, where the mask contained zeros at the positions corresponding to higher frequencies. The purpose of this step is to remove steganography from the generated images, as the steganography signal, introduced into the images by the generators to ease up reconstruction, is of low amplitude and high frequency. Specifically, we performed low-pass filtering by nullifying the coefficients corresponding to the 50% highest frequencies. To obtain an image $\hat{\mathbf{y}}^*$ or $\hat{\mathbf{x}}^*$ from



(a) Generated image (b) Filtered image (c) Difference image

Figure 4: Generated images (a) filtered with our proposed DCT filtering are almost indistinguishable since the filter is designed only to remove steganography. In (c) the absolute difference between the filtered and the generated image is shown.

the filtered coefficients C^* , the inverse DCT is applied by computing:

$$\text{iDCT}(i, j) = \sum_{i=0}^{N-1} \sum_{j=0}^{M-1} \alpha(i)\alpha(j)C^*(i, j)hv. \quad (5)$$

Naturally, the filtering does not only affect the hidden signal but any high-frequency image information, such as foci inside of the nuclei. However, we argue that for reconstructing masks from nuclei images, these high-frequency features are not mandatory, and the differences are subtle. A side-by-side comparison between an original image and a DCT filtered image and their absolute difference is depicted in Figure 4c. As the cycle loss mainly promotes the CycleGAN steganography, we apply the filtering only to images used for computation of the cycle loss, *i.e.*, the adversarial losses are computed on the unfiltered generator output. Figure 3 gives an overview of the CycleGAN data flow with our proposed DCT filtering for synthetic dataset generation.

2.3. Mask synthesis

Böhlend *et al.* show that ellipses are a simple but suitable model for representing the shape of nuclei [1] and can be used to create synthetic label images. However, careful parameterization of the ellipses is required to match the properties of the nuclei of real microscopy images. A substantial deviation between synthetic and real distribution might result in CycleGAN being unable to learn the mapping correctly. The standard equation for an ellipse is

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1, \quad (6)$$

with the semi-major axis a and the semi-minor axis b . To express how elongated an ellipse is, we use the eccentricity e , which is for $a > b$ defined as:

$$e = \sqrt{1 - \left(\frac{b}{a}\right)^2}. \quad (7)$$

For mask synthesis, we randomly sample the eccentricity in the range $[0.4, 0.9]$ to compute the minor axis b with:

$$b = \sqrt{1 - e^2}a. \quad (8)$$

The size of the semi-major axis a is randomly sampled from a uniform range which is determined on a per-dataset basis. The uniform range for the number of nuclei per image is estimated on sample images from each dataset. A random rotation with a random angle from a uniform distribution in the range $[0^\circ, 179^\circ]$ is applied to each nucleus individually. Nuclei are placed in the synthetic mask images so that they may touch but never overlap.

3. Experiments

3.1. Evaluation setup

To evaluate the influence of DCT filtering on CycleGAN training, we generate synthetic datasets using a CycleGAN with DCT filtering and an unmodified CycleGAN as a baseline. Additionally, we implement two methods from literature: the denoising approach (*DnCNN*) from [7] and the noise-injection approach (*Noise*) using Gaussian noise as proposed in [5]. Instead of using the synthetic input mask from our mask synthesis model and the generated microscopy image by CycleGAN for our synthetic dataset, we use the cycled input mask along with the generated microscopy image. This is motivated by the fact that our method aims to increase correspondence between generated image and cycled mask by filtering the generated image. Additionally, ellipses are a suitable but imperfect approximation of the nuclei shape. If the CycleGAN deviates slightly from this approximation, the cycled mask allows adaptation to these deviations. This is under the assumption that CycleGAN steganography was removed successfully. However, the cycled masks generated by CycleGAN are close to binary masks that might contain minor imperfections. To convert the binary masks into instance masks, we close small holes in the masks and compute the Euclidean distance transform of the binary mask and then apply Watershed segmentation with morphological operators. An example of this post-processing step is visualized in Figure 5. The synthetic dataset is then utilized for training the nuclei segmentation network StarDist [8]. Subsequently, the trained segmentation network is applied to real microscopy data for testing. We use StarDist’s matching function to compute precision and recall with IoU thresholds of $\tau = 0.5$ and $\tau = 0.75$. Additionally, we report the F1 score. Each dataset is generated three times with random initialization and results are averaged.

3.2. Datasets

To evaluate our method, we train CycleGAN on two public FM nuclei datasets from the Broad Bioimage Benchmark

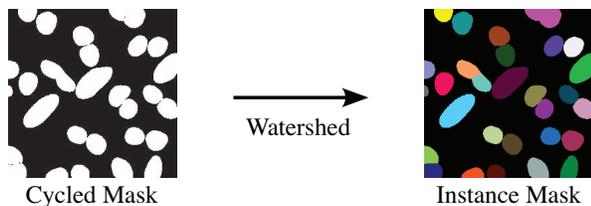


Figure 5: To convert the binary cycled masks into instance masks that can be used to train segmentation networks, we apply the watershed algorithm on the distance-transformed cycled masks.

Image Collection [2]. For each nuclei image, we normalize pixel values into the range $[-1, 1]$ and extract random crops of 256×256 pixels. To create synthetic labels, we use the mask synthesis model described in Section 2.3. The parameters for mask synthesis can be obtained from Table 1 for the DSB dataset and Table 2 for dataset BBBC039v1. They were determined based on a small set of three selected nuclei images for each dataset individually. During the training of CycleGAN, new random masks are created for each batch. After the training of CycleGAN, a fixed set of 1500 synthetic masks is used to generate nuclei images as a training set for the downstream nuclei segmentation network.

3.2.1 Data Science Bowl (DSB) 2018 Nuclei Dataset

The Data Science Bowl (DSB) 2018 Nuclei Dataset was released for the 2018 Data Science Bowl Kaggle challenge and is listed in the Broad Bioimage Benchmark Collection as BBBC038v1. The dataset contains thousands of nuclei obtained from various organisms, including humans, mice, and flies. Nuclei in the dataset appear in different contexts and states, *e.g.*, cell division, genotoxic stress, and differentiation. While the original dataset contains a variety of treatments (*e.g.*, hematoxylin and eosin stained nuclei), only FM images were used in this work. The dataset contains two labeled subsets, *stage1_train* with 538 FM images and *stage1_test* with 53 FM images. Additionally, the dataset contains an unlabelled subset, *stage2_test* of which we select 667 FM images. We use *stage1_train* and *stage2_test* for CycleGAN training. The downstream nuclei segmentation networks were tested on *stage1_test*. Figure 6 visualizes the variety of the dataset using three example images.

3.2.2 BBBC039v1 dataset

The dataset BBBC039v1 is a FM dataset, and the images were obtained as part of a high-throughput chemical screen on U-2 OS cells by Caicedo *et al.* [3]. The dataset presents a variety of nuclei shapes, different degrees of clustering, and consists of a total of 230000 single nuclei, which were man-

Table 1: Parameters of the mask synthesis model of DSB dataset. Each synthetic mask’s parameters are chosen from random uniform intervals denoted below. By using multiple subsets of parameter combinations, we allow the mask synthesis model to reflect the diverse characteristics of the DSB dataset.

Length major axis	Nuclei count	Eccentricity e
(5, 10)	(1, 150)	(0.4, 0.9)
(10, 15)	(1, 40)	(0.4, 0.9)
(15, 20)	(1, 40)	(0.4, 0.9)
(20, 25)	(1, 40)	(0.4, 0.9)
(25, 30)	(1, 20)	(0.4, 0.9)
(30, 35)	(1, 20)	(0.4, 0.9)

ually annotated. However, compared to the DSB dataset, it exhibits much less variability in shape and nuclei style, as all images were obtained from the same experiment. The dataset holds 200 unique images, each of size 520×696 pixels. Of these, 150 were used for CycleGAN training. The remaining 50 images were used as the test set for the downstream nuclei segmentation network. Exemplary images are shown in Figure 7.

Table 2: Parameters of the mask synthesis model of BBBC039v1 dataset. Each synthetic mask’s parameters are chosen from random uniform intervals denoted below. Compared to DSB, the BBBC039v1 dataset is less diverse concerning nuclei shape and distribution, reflected by our selected parameters.

Length major axis	Nuclei count	Eccentricity e
(10, 20)	(20, 60)	(0.6, 0.9)
(20, 40)	(20, 30)	(0.6, 0.9)

3.3. Network training

3.3.1 CycleGAN

All compared CycleGAN variants are based on the network architecture described in the original CycleGAN paper [13] and we used the official PyTorch implementation.¹ We trained all networks for 100 epochs with a learning rate of $2e-4$, which was linearly decayed to 0 for another 100 epochs, using an Adam optimizer ($\beta_1 = 0.5$ and $\beta_2 = 0.999$) with a batch size of 1. To stabilize the training, the discriminator was updated from an image pool of 50 images [9]. The weight λ for the cycle-consistency loss was set to $\lambda_F = 15.0$ for generator F and to $\lambda_G = 10.0$ for generator G . We set λ_G lower since the mapping from masks to nuclei images is not unique, as it is not directly possible to

¹Available at <https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix/>.

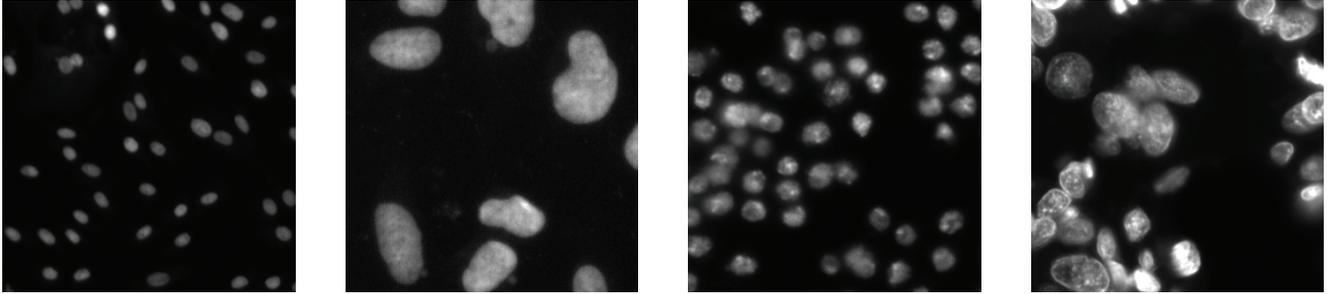


Figure 6: Sample images from DSB dataset.

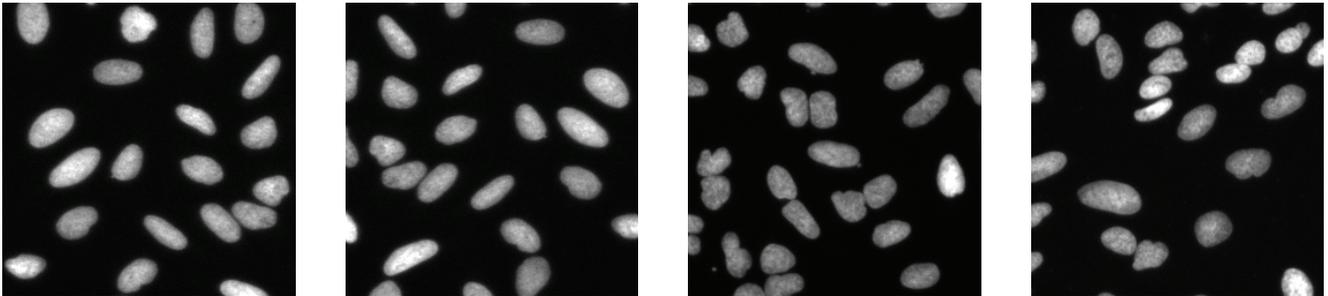


Figure 7: Sample images from BBBC039v1 dataset.

infer the texture from a nucleus based on a binary mask. For our comparison with [7], we use the reference implementation of *DnCNN* [12], which is trained for blind Gaussian denoising with a large range of noise levels ($\sigma \in [0, 55]$). The denoising network weights were frozen during training and inference runs. Denoising is only applied for computation of the cycle loss.

3.3.2 StarDist

We use StarDist [8] as downstream nuclei segmentation network. It is a popular choice in the field due to its effectiveness, versatility, and state-of-the-art nuclei segmentation performance. However, in this work, the focus is not on achieving the best possible result but on investigating the applicability of differently generated synthetic datasets for a downstream task. We train StarDist for 50 epochs with a batch size of 16 and a learning rate of 3×10^{-4} . We increased the number of rays used for the star-convex polygon representation from the default 32 to 64 in order to obtain more detailed polygon shapes. All remaining parameters of StarDist were left unchanged.

4. Results

4.1. Assessment of image fidelity

While the primary goal of our proposed methods is to improve the downstream task of nuclei segmentation, we report the results of the fidelity comparison for generated

and original images in Table 3 for completeness. To this end, we compute the Fréchet inception distance (FID) as proposed in [6]. The CycleGAN with DCT filtering yields the best FID score on the DSB dataset (96.4) and second-lowest on BBBC039v1 (328.1). The low standard deviation (± 0.7 for the three repetitions with DSB dataset) indicates that DCT filtering can also help to produce more stable results. We generally observe unusually high FID scores among all datasets. However, the FID score computation is based on a feature extractor trained on natural images. This representation might not be optimal for quality assessment of FM images, leading to higher scores and limited validity.

Table 3: FID score (lower is better) is chosen as metric to compare the image fidelity of generated data for datasets DSB and BBBC039v1 with original images from the respective datasets

Model	DSB	BBBC039v1
Orig. CycleGAN	110.1 ± 3.0	380.2 ± 12.9
w. DCT (ours)	96.4 ± 0.7	328.8 ± 13.8
w. noise [5]	104.9 ± 4.6	376.6 ± 39.7
w. DnCNN [7]	101.3 ± 3.5	327.1 ± 4

4.2. Qualitative results of synthetic images

A representative sample of generated images using the discussed methods is depicted along with (from left to right)

the generating input mask, the respective cycled mask, and the post-processed instance mask as an overlay in Figure 8. For all four analyzed methods, we observe a deviation from the original input mask. This deviation is not necessarily negative as long as it is represented in the cycled mask. This is not the case for the unmodified CycleGAN. Multiple nuclei instances are not rendered into the image but occur in the cycled mask. In this case, it is almost a pixel-perfect copy of the original input mask. The reason for this is the before-mentioned CycleGAN steganography. If our proposed DCT filtering is used during training, we observe no steganography during test time. The cycled mask resembles the generated image closely. We also observe that the nuclei do not exhibit an exact ellipse shape, which we consider a positive effect. This allows the model to adapt to the shape of nuclei found in the training set and overcome shortcomings of the ellipse synthesis representations. For the approaches from literature, specifically the filtering with the denoising network *DnCNN* [7], we observe that steganography was reduced but not fully removed as there are nuclei in the cycled mask that were not present in the generated image. For the noise injection approach from [5], we observe the largest deviation from the input mask in the generated image. However, since this is reflected in the cycled mask, we conclude that removing the steganography was successful. We also observe a slightly less realistic reproduction of the typical nuclei texture found in the remaining images when noise injection is used.

4.3. Segmentation performance on DSB

The results for the segmentation network tested on the real-data test split *stage1_test* of the DSB dataset are reported in Table 4 for an IoU threshold of $\tau = 0.5$ and in Table 5 for an IoU threshold $\tau = 0.75$. The best results in terms of precision are achieved with the training set created by our proposed CycleGAN with DCT filtering, outperforming the baseline CycleGAN by 19.7 percentage points for $\tau = 0.5$ and 38.7 percentage points for $\tau = 0.75$. For the approach using *DnCNN* we observe comparable performance at $\tau = 0.5$, with a larger gap (8.3 percentage points less precision compared to DCT) for the more challenging threshold of $\tau = 0.75$. The noise-injection approach suffers from a high variance leading to a slightly worse average performance.

4.4. Segmentation performance on BBC039v1

We report results for the downstream segmentation task on the BBC039v1 test set for $\tau = 0.5$ and $\tau = 0.75$ in Table 6 and Table 7 respectively. We observe a high precision for all methods except the noise injection, indicating a low number of false positives and a high number of true positives. Compared to the DSB dataset, we observe higher scores among all methods and metrics due to the

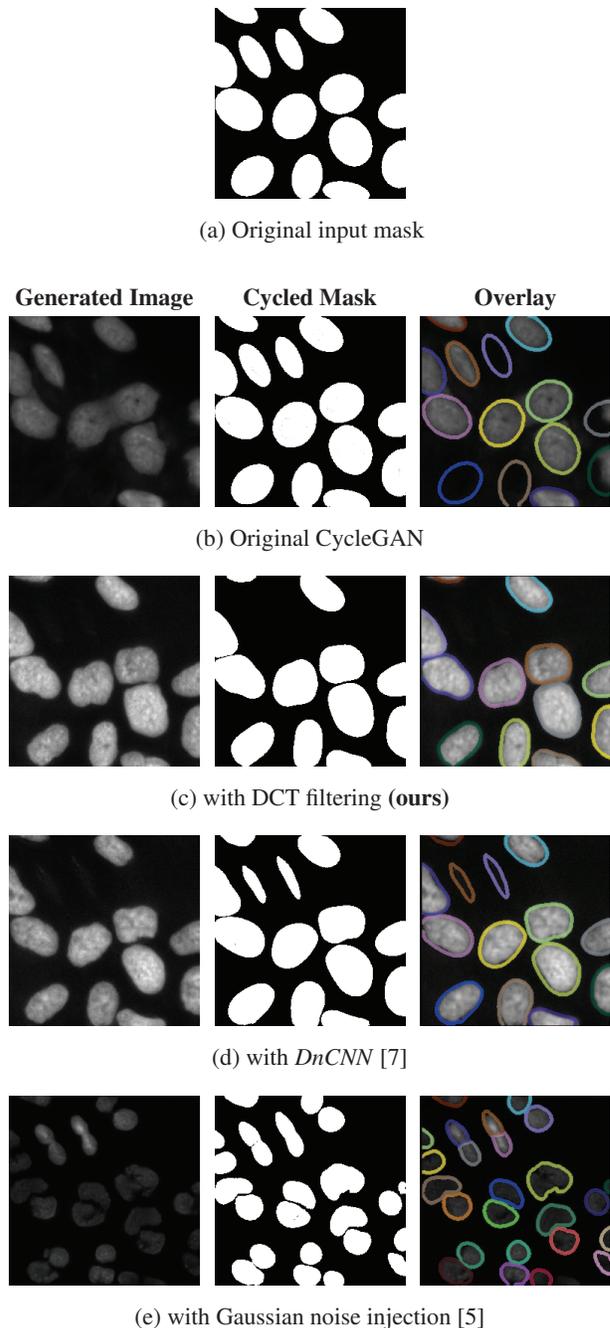


Figure 8: Generated images based on the DSB dataset along with the cycled mask. The final segmentation mask used for StarDist training is shown in the third column as an overlay over the generated image.

lower diversity and therefore complexity of the data set. The lower complexity results in a lower amount of steganography. Consequently, the methods differ only slightly in their F1-score, with DCT filtering being the best method by a margin of 2 percentage points ($\tau = 0.5$) compared to the

Table 4: StarDist segmentation performance in % on DSB *stage1_test* split, trained on synthetic datasets generated from different CycleGAN methods. Scores are computed for $\tau = 0.5$

Model	Precision	Recall	F1-Score
Orig. CycleGAN	74.3 \pm 5	63.7 \pm 3	68.6 \pm 4
w. DCT (ours)	94.0 \pm 1	61.1 \pm 1	74.0 \pm 1
w. noise [5]	87.0 \pm 14	39.8 \pm 31	50.3 \pm 37
w. DnCNN [7]	92.3 \pm 1	62.8 \pm 3	74.7 \pm 2

Table 5: StarDist segmentation performance in % on DSB *stage1_test*, trained on synthetic datasets generated from different CycleGAN methods. Scores are computed for $\tau = 0.75$, which requires more precise predictions compared to $\tau = 0.5$

Model	Precision	Recall	F1-Score
Orig. CycleGAN	49.9 \pm 5	40.2 \pm 3	43.3 \pm 4
w. DCT (ours)	85.6 \pm 1	55.6 \pm 1	67.4 \pm 1
w. noise [5]	67.7 \pm 32	35.5 \pm 29	44.5 \pm 36
w. DnCNN [7]	77.3 \pm 11	52.4 \pm 5	62.4 \pm 7

baseline. The margin between the proposed DCT filtering and the baseline CycleGAN methods increases for a higher threshold ($\tau = 0.75$) to 8.4 percentage points in terms of F1-score. This indicates that DCT filtering leads to more realistic nuclei shapes. Similar to the results of the DSB dataset, we observe a high variability in performance for the noise-injection results. The reason of this is likely due to noise injection causing the GAN training to be less stable, generating more deficient samples.

Table 6: StarDist segmentation performance in % on BBBC039v1 test set, trained on synthetic datasets generated from different CycleGAN methods. Scores are computed for $\tau = 0.5$

Model	Precision	Recall	F1-Score
Orig. CycleGAN	99.9 \pm 1	86.3 \pm 6	92.1 \pm 4
w. DCT (ours)	97.9 \pm 1	90.7 \pm 1	94.1 \pm 1
w. noise [5]	82.1 \pm 29	76.3 \pm 28	79.1 \pm 28
w. DnCNN [7]	99.9 \pm 1	82.4 \pm 2	90.3 \pm 2

5. Discussion

In this work, we proposed a new anti-steganography approach for CycleGANs and evaluated existing methods to improve CycleGAN consistency for FM nuclei generation. We further showed the performance of nuclei segmentation trained solely on synthetic data. Using our proposed method, we observe a qualitative improvement in the

Table 7: StarDist segmentation performance on BBBC039v1 test set, trained on synthetic datasets generated from different CycleGAN methods. Scores are computed for $\tau = 0.75$

Model	Precision	Recall	F1-Score
Orig. CycleGAN	82.8 \pm 11	72.6 \pm 15	77.3 \pm 13
w. DCT (ours)	89.2 \pm 10	82.5 \pm 8	85.7 \pm 9
w. noise [5]	65.2 \pm 54	61.1 \pm 50	61.1 \pm 52
w. DnCNN [7]	91.4 \pm 11	75.5 \pm 10	82.6 \pm 10

generated images and masks and outperform existing anti-steganography approaches quantitatively on a downstream segmentation task. A considerable advantage of DCT filtering compared to using a *DnCNN* is the reduced computational overhead. Instead of feeding the generated images through a neural network (*i.e.*, the *DnCNN*), we make use of the efficient and computationally cheap implementation of the DCT algorithm. By pairing the cycled mask with the generated image for synthetic dataset creation, we allow the network to deviate from the input mask, allowing CycleGAN to overcome potential shortcomings of the mask synthesis model. This can be a crucial factor in enabling CycleGAN to learn the desired mapping. In [1] Böhlend *et al.* show that CycleGAN can only produce satisfactory results if the shape and occurrence of synthetic masks are similar to the ground-truth distribution. In unsupervised settings, the ground-truth distribution is not available, and estimating the parameters of the mask synthesis model relies on a limited number of samples, which can introduce challenges and uncertainties. A limitation of using the cycled masks is the necessity of introducing post-processing steps to convert them into instance masks. This in turn might induce an unknown bias into the processed masks and is thus prone to errors. This could be avoided by directly generating instance masks in CycleGAN, *e.g.*, by rendering additional boundaries of masks or using a StarDist-like mask formulation based on star-convex polygons.

In this work, we did not consider the segmentation capabilities of CycleGAN itself. Due to its architectural design, CycleGAN is not only trained to generate microscopy images from masks but also to segment microscopy images into masks. In future work, we want to explore to what extent this allows this allows CycleGAN to be used as an unsupervised nuclei segmentation network, *e.g.*, by using segmentation losses such as the Dice loss for cycle consistency. Exploring the generalizability of the proposed steganography removal approach to other cycle-loss-based image-to-image translation tasks could open up new possibilities for improving various GAN-based applications.

6. Acknowledgments

We gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Projektnummer 405969122 (J. U., S. U., K. B.) and Dhip campus - Bavarian aim (K. B., M. S., T. W.) M. T., F. W., M. G., and S. U. received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC Grant No. 810316).

References

- [1] M. Böhlend, T. Scherr, A. Bartschat, R. Mikut, and M. Reichl. Influence of synthetic label image object properties on GAN supported segmentation pipelines. In *Proceedings - 29. Workshop Computational Intelligence, Dortmund, 28. - 29. November 2019. Ed.: F. Hoffmann, E. Hüllermeier, R. Mikut*, pages 289–309. KIT Scientific Publishing, 2019. 47.01.02; LK 01.
- [2] Juan C. Caicedo, Allen Goodman, Kyle W. Karhohs, Beth A. Cimini, Jeanelle Ackerman, Marzieh Haghighi, CherKeng Heng, Tim Becker, Minh Doan, Claire McQuin, Mohammad Rohban, Shantanu Singh, and Anne E. Carpenter. Nucleus segmentation across imaging experiments: the 2018 data science bowl. *Nature Methods*, 16(12):1247–1253, Oct. 2019.
- [3] Juan C. Caicedo, Jonathan Roth, Allen Goodman, Tim Becker, Kyle W. Karhohs, Matthieu Broisin, Csaba Molnar, Claire McQuin, Shantanu Singh, Fabian J. Theis, and Anne E. Carpenter. Evaluation of deep learning strategies for nucleus segmentation in fluorescence images. *Cytometry Part A*, 95(9):952–965, 2019.
- [4] Casey Chu, Andrey Zhmoginov, and Mark Sandler. CycleGAN, a master of steganography. *arXiv preprint arXiv:1712.02950*, 2017.
- [5] Jonas Löhdefink and Tim Fingscheidt. Improving performance of semantic segmentation cyclegans by noise injection into the latent segmentation space. *arXiv preprint arXiv:2201.06415*, 2022.
- [6] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, 2022.
- [7] Horia Porav, Valentina Musat, and Paul Newman. Reducing steganography in cycle-consistency GANs. In *CVPR Workshops*, pages 78–82, 2019.
- [8] Uwe Schmidt, Martin Weigert, Coleman Broaddus, and Gene Myers. Cell detection with star-convex polygons. In Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pages 265–273, Cham, 2018. Springer International Publishing.
- [9] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb. Learning from simulated and unsupervised images through adversarial training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2107–2116, 2017.
- [10] Carsen Stringer, Tim Wang, Michalis Michaelos, and Marius Pachitariu. Cellpose: a generalist algorithm for cellular segmentation. *Nature Methods*, 18(1):100–106, 2021.
- [11] Gregory K Wallace. The jpeg still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992.
- [12] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017.
- [13] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251. IEEE, 2017.