Stochastic Digital Twin for Copy Detection Patterns

Yury Belousov, Olga Taran, Vitaliy Kinakh and Slava Voloshynovskiy

Department of Computer Science, University of Geneva, Switzerland

{yury.belousov, olga.taran, vitaliy.kinakh, svolos}@unige.ch

Abstract—Copy detection patterns (CDP) present an efficient technique for product protection against counterfeiting. However, the complexity of studying CDP production variability often results in time-consuming and costly procedures, limiting CDP scalability. Recent advancements in computer modelling, notably the concept of a "digital twin" for printing-imaging channels, allow for enhanced scalability and the optimization of authentication systems. Yet, the development of an accurate digital twin is far from trivial.

This paper extends previous research which modelled a printing-imaging channel using a machine learning-based digital twin for CDP. This model, built upon an information-theoretic framework known as "Turbo", demonstrated superior performance over traditional generative models such as CycleGAN and pix2pix. However, the emerging field of Denoising Diffusion Probabilistic Models (DDPM) presents a potential advancement in generative models due to its ability to stochastically model the inherent randomness of the printing-imaging process, and its impressive performance in image-to-image translation tasks.

This study aims at comparing the capabilities of the Turbo framework and DDPM on the same CDP datasets, with the goal of establishing the real-world benefits of DDPM models for digital twin applications in CDP security. Furthermore, the paper seeks to evaluate the generative potential of the studied models in the context of mobile phone data acquisition. Despite the increased complexity of DDPM methods when compared to traditional approaches, our study highlights their advantages and explores their potential for future applications.

Index Terms—Copy detection patterns, machine learning, digital twin, denoising diffusion model, TURBO, CycleGAN, pix2pix.

I. INTRODUCTION

The recent upsurge in the utilization of Copy Detection Patterns (CDP), as described in [1]–[4], has emerged as a viable method for safeguarding products against counterfeiting practices. However, the exploration of variability inherent in CDP production represents a process that is both timeintensive and financially demanding. This process necessitates the acquisition of vast volumes of data, a requirement that places a significant constraint on the scalability of the approach to incorporate new products, manufacturing technologies, and imaging devices. Consequently, the expansive adoption and continued research into CDP are impeded.

To overcome these limitations and promote the ongoing advancement of CDP, an approach involving computational modeling of the entire production pipeline, incorporating the printing and imaging channels, has been proposed. This method is referred to as a *digital twin* [5]. Through this approach, a comprehensive and accurate simulation of the production process is generated, allowing for an efficient examination and optimization of CDP without the traditional restrictions associated with physical data acquisition. This



Fig. 1. Schematic block-diagram of investigated DDPM generative model g_{φ} . The DDPM generative model can generate K synthetic CDP $\{\tilde{\mathbf{x}}_i^k\}_{k=1}^K$ from digital template \mathbf{z} and vice versa generate K synthetic templates $\{\tilde{\mathbf{z}}_i^k\}_{k=1}^K$ for a given \mathbf{x} . We show only the first case. The stochasticity of the generative process is ensured by different noise realizations $\boldsymbol{\epsilon}$.

approach offers considerable potential for increasing the efficiency and effectiveness of anti-counterfeiting measures based on CDP.

The design of *digital twin* for printing-imaging channels is not a trivial task but it is crucial for both the defender and attacker. If successful, it will enable the overall optimization of the whole authentication system and, in particular, the optimization of the estimation of digital templates from the physical samples and synthesis of CDP images from the corresponding digital templates. Moreover, it simplifies the modeling of the intra-class variabilities and the investigation of adversarial examples. The number of training pairs needed for *digital twin* is small (in the order of hundreds), while the trained model can be applied to millions of unseen digital templates.

The current work is a continuation of our previous work [5] that was dedicated to modeling a printing-imaging channel using a machine learning-based *digital twin* for CDP. The model studied in [5] is based on an information-theoretic framework called *Turbo*. In our current work, we aim at comparing Turbo to the Denoising Diffusion Probabilistic Models (DDPM) [6], which present a popular family of modern generative models (Fig. 1). DDPM model the prior data distribution via a diffusion process. Recently the DDPM methods demonstrated remarkable performance in the image-to-image translation tasks and outperformed many state-of-the-art models based on GAN-like architectures [7]. In contrast to the GAN-based generators that are mostly deterministic in nature, DDPM allows stochastic outputs, i.e., the different



Fig. 2. The stochasticity in the DDPM Model. The first column displays the original digital template \mathbf{z} at the top and its counterpart physical sample \mathbf{x} at the bottom. The subsequent columns show the stochastic estimations of the CDP images $\{\tilde{\mathbf{x}}^k\}_{k=1}^{5}$ and the digital templates $\{\tilde{\mathbf{z}}^k\}_{k=1}^{5}$, generated by the DDPM based on the Palette framework. To enhance visual comprehension, only an enlarged 11×11 central crop is displayed.

outputs for the same input. Taking into account the natural randomness of the printing-imaging process, the stochasticity of the synthesised twins is a key factor for high-precision simulation of real CDP. Besides this valuable advantage, the DDPM methods have high complexity compared to the traditional approaches. That is why the study of real advantages of DDPM based CDP digital twins represents a great practical interest.

In our previous work [5] we demonstrated the superiority of the Turbo framework over the state-of-the-art generative models. The main goal of this study is to compare Turbo with DDPM on the same CDP datasets and to establish the real advantages of DDPM models. Moreover, we aim at evaluating the generative capabilities of the models in the context of mobile phone data acquisition.

II. RELATED WORK

A. Turbo family

The Turbo framework was derived based on the solid information theoretic foundations [5]. That framework consists of two paths, i.e., direct and reverse ones. In the general case, both paths are trained simultaneously and share common training blocks. However, in particular cases, the framework might be trained in one path only. At the inference stage, both paths or just one of them might be used depending on the targeted application. As it was investigated in [5], Turbo can be trained on paired or unpaired data, providing flexibility in its application. Turbo generalizes pix2pix (paired) [8] and CycleGAN (unpaired) [9] image-to-image translation systems.

Turbo consists of several building blocks and losses that make the training procedure complex enough. Once trained, Turbo is quite efficient and very fast at the inference stage. The main drawback of Turbo is the deterministic nature of generated CDP, i.e., for the given input it provides only one output.

In [5] we extensively investigated the impact of various factors such as the backbone architectures, the discriminator types, losses, etc., on the overall system's performance and

found the optimal ones. In the current work, we use two found optimal configurations:

• TURBO^{paired (w \mathcal{D}) • \mathcal{T} URBO^{paired (w \mathcal{D}) $\mathcal{L}^{\text{paired (w }\mathcal{D})}_{\text{CNN-RESNET-CNN}}(\phi, \theta) = \mathcal{L}_{\tilde{z}}(\mathbf{z}, \tilde{\mathbf{z}}) + \mathcal{D}_{\tilde{z}}(\mathbf{z}, \tilde{\mathbf{z}})$ $+ \lambda_D \mathcal{L}_{\hat{x}}(\mathbf{x}, \hat{\mathbf{x}}) + \lambda_D \mathcal{D}_{\hat{x}}(\mathbf{x}, \hat{\mathbf{x}})$ $+ \lambda_T \mathcal{L}_{\hat{x}}(\mathbf{x}, \tilde{\mathbf{x}}) + \lambda_T \mathcal{D}_{\hat{x}}(\mathbf{x}, \tilde{\mathbf{x}})$ $+ \lambda_T \lambda_R \mathcal{L}_{\hat{z}}(\mathbf{z}, \hat{\mathbf{z}}) + \lambda_T \lambda_R \mathcal{D}_{\hat{z}}(\mathbf{z}, \hat{\mathbf{z}}),$ • TURBO^{paired (w/o \mathcal{D})}}}

$$\mathcal{L}_{\text{UNET}}^{\text{paired (w/o D)}}(\phi, \theta) = \mathcal{L}_{\tilde{\mathbf{z}}}(\mathbf{z}, \tilde{\mathbf{z}}) + \lambda_D \mathcal{L}_{\hat{\mathbf{x}}}(\mathbf{x}, \hat{\mathbf{x}}) + \lambda_T \mathcal{L}_{\tilde{\mathbf{x}}}(\mathbf{x}, \tilde{\mathbf{x}}) + \lambda_T \lambda_R \mathcal{L}_{\hat{\mathbf{z}}}(\mathbf{z}, \hat{\mathbf{z}}),$$

where **x** denotes the image of CDP, **z** denotes the digital template, $\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$ denote the reconstructions and $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{z}}$ are the generated images. The terms $\mathcal{L}_{\tilde{\mathbf{z}}}(\mathbf{z}, \tilde{\mathbf{z}})$, $\mathcal{L}_{\hat{\mathbf{x}}}(\mathbf{x}, \hat{\mathbf{x}})$, $\mathcal{L}_{\tilde{\mathbf{x}}}(\mathbf{x}, \tilde{\mathbf{x}})$ and $\mathcal{L}_{\hat{\mathbf{z}}}(\mathbf{z}, \hat{\mathbf{z}})$ are the conditional cross-entropy terms that are implemented as ℓ_1 -norm pair-wise losses between the corresponding entities. $\mathcal{D}_{\tilde{\mathbf{z}}}(\mathbf{z}, \tilde{\mathbf{z}})$, $\mathcal{D}_{\hat{\mathbf{x}}}(\mathbf{x}, \hat{\mathbf{x}})$, $\mathcal{D}_{\tilde{\mathbf{x}}}(\mathbf{x}, \tilde{\mathbf{x}})$ and $\mathcal{D}_{\hat{\mathbf{z}}}(\mathbf{z}, \hat{\mathbf{z}})$ impose Kullback-Leibler (KL)-divergence constraints, i.e., the distribution matching losses, a.k.a. adversarial losses, between the corresponding distributions and the parameters λ_T , λ_D and λ_R trade-off the losses. The detailed development of Turbo's losses and the schematic representation of the direct and reverse paths are given in [5] and the corresponding code.

B. DDPM

The DDPM [6] are based on the minimization of Fisher divergence, which is also closely linked with the KL-divergence, between the data distribution and the energy-based model approximating the data distribution. The core concept of DDPM is to use a score function representing the gradient of the logarithm of the energy-based model with respect to the data sample to suppress the dependence on the normalization constant, which is infeasible to compute in practice [10]. Similar to Turbo, DDPM consists of forward and reverse paths. However, in contrast to Turbo, the DDPM forward path is not trainable and is based on the addition of noise to network input. The addition of noise with variable variance



Fig. 3. An example of 2D variability for a randomly selected CDP. We used pixel-wise standard deviation σ to estimate the variability among the generated images. For better visual comprehension, we display a central crop that is equal to half the dimensions of the full image.

aims at "interpolation" of data distribution represented by sparse training data samples. The variable variance of noise should address the different regions of data distribution in a function of the estimated probability density function [10].

From the point of view of the CDP nature, the trained Turbo model can produce the printing simulations and the digital template estimations simultaneously, while DDPM requires training of two separate models. With respect to the number of optimised losses, the DDPM training is simpler and includes only one loss. We use the Palette model [11] to implement conditional DDPM. For the $z \rightarrow \tilde{x}$ case, the model's loss is:

$$\mathcal{L}^{DDPM}(\varphi) = \mathbb{E}_{t,\mathbf{z},\mathbf{x},\boldsymbol{\epsilon}} \left[\left\| \boldsymbol{\epsilon} - g_{\varphi} \left(\sqrt{\bar{\alpha}_{t}} \mathbf{x} + \sqrt{1 - \bar{\alpha}_{t}} \boldsymbol{\epsilon}, \mathbf{z}, t \right) \right\|^{2} \right]$$

where x denotes the target image, z denotes the digital template used as a conditioning, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ denotes the noise added at step t, g_{φ} stands for the parametrized denoiser model, $\bar{\alpha}_t$ denotes the noise scale parameter [6]. For the $\mathbf{x} \rightarrow \tilde{\mathbf{z}}$ channel modeling, a similar loss is used, but the digital template z is used as the target image, and the model is conditioned by x. In contrast to Turbo, the DDPM loss does not allow the training on unpaired data.

Training and inference stages are iterative and require adapting to many noise levels. Contrary to Turbo, which generates data in a single step, DDPM might need hundreds of steps to produce the final result. Despite this, a notable distinction between DDPM and Turbo lies in the stochastic nature of DDPM, enabling it to generate multiple outputs from a single input, thereby accommodating the intrinsic randomness associated with the printing process. The schematic block diagram of DDPM is shown in Fig. 1.

III. DATASET AND TRAINING DETAILS

A. Dataset

For empirical evaluation of the models under investigation we used the data acquired by two modern mobile phones and by a high-resolution scanner.

The experiments on the scanner data are an extension of our previous work [5]. In this respect, the same Indigo 1×1 symbol dataset $[12]^1$ was used. This dataset consists of 720 digital templates of size 228×228 with 1×1 pixel symbol size. The digital templates have been printed at HP Indigo 7600 industrial printer at a resolution of 812.8 dpi and enrolled by Epson Perfection V850 Pro scanner at a resolution of 2400 ppi. Considering the ratio between the printing and acquisition resolutions, the obtained CDP are of size 684×684 , i,e., 1×1 pixel in the digital template corresponds to a 3×3 block in the acquired CDP. The final codes are 16-bit grayscale images.

¹http://sip.unige.ch/projects/snf-it-dis/datasets/indigo-base



Fig. 4. The x-axis represents the 512 different possible patterns ω ordered by their flattened binary representations. The y-axis represents the standard deviation of the central pixel of each pattern for the iPhone.

The experiments on the mobile phone data were performed on the recently created Indigo 1×1 variability dataset $[13]^2$ that consists of 1440 digital templates of size 228×228 with 1×1 pixel symbol size. The templates have been printed at HP Indigo 5500 industrial printer at a resolution of 812.8 dpi and enrolled by iPhone 12 Pro and Samsung Galaxy Note 20 Ultra cell phones. The obtained CDP images are of size 228×228 and encoded as 8-bit RGB images. However, for the sake of simplicity, we convert them into grayscale images.

Both mentioned datasets contain original and fake CDP. For our experiments, we used only the original codes. However, it should be noted that in both cases the fakes were produced on the same printing and acquisition equipment as the original codes. In this respect, the model trained on the original codes can be effectively applied to generate fake codes.

B. Training details

The Turbo framework architectures' details and training conditions are the same as in [5].

As a DDPM-based framework, we used the Palette model [11] with the UNET architecture inspired by [7]. We modify UNET in the following way: it takes two channels as input, where the first channel is a noise and the second one is used for conditioning; the model incorporates attention with resolutions of 16 and includes two residual blocks per downsampling step. We initialized the model using Kaiming initialization. Additionally, we set the dropout rate to 0.2 to prevent overfitting due to the high similarity between the CDP. We train our models on a single A100 GPU with 80GB of memory with a mini-batch of size 36 for 15000 training epochs. We use a standard Adam optimizer with the $5e^{-5}$ learning rate and without a learning rate warmup schedule. Similarly to [11], we use 0.9999 EMA but, during the inference, we do not perform the hyper-parameter tuning over noise schedules and refinement steps. During training, we employ the same linear





Fig. 5. The x-axis represents the 512 different possible patterns ω ordered by their flattened binary representations. The y-axis represents the probability of bit-flipping for the central pixel of each pattern computed from iPhone data.

noise schedule of $(1e^{-6}, 0.01)$ with 2000 time-steps and 1000 refinement steps with a linear schedule of $(1e^{-4}, 0.09)$ during inference as in [11]³.

IV. RESULTS AND DISCUSSION

A. Palette model stochasticity

Fig. 2 demonstrates several examples of the diverse outputs of the Palette model produced for the same randomly selected input from the iPhone subset. The top row corresponds to the modeling of printing channel $z \to \tilde{x}$ and the bottom one shows the digital template estimation channel $x \to \tilde{z}$.

To illustrate the variability in the generated data, we picked the same template and stacked the produced outputs as a 3D tensor, then we calculated the standard deviation in the image dimension, i.e., for each pixel of generated images. The obtained results are visualized in Fig. 3. The left part of the figure shows the results for the $\mathbf{x} \to \tilde{\mathbf{z}}$ for the iPhone, Samsung, and scanner, respectively. It is straightforward to observe numerous regions characterized by diminutive standard deviation, as indicated by the dark blue hue. These regions symbolize the model's degree of confidence in the generated outcomes, which correlates to the conglomerations of white and black pixels. Conversely, the yellow hue denotes areas of heightened standard deviation, reflective of the model's uncertainty. These areas typically align with the transitional regions, manifesting the boundary conditions between different pixel clusters. It is important to note that these features are best seen in the case of the scanner due to the higher acquisition resolution. The right part shows the results for the $z \to \tilde{x}$ channel. In contrast to the $\mathbf{x} \to \tilde{\mathbf{z}}$ channel, the general dynamic range for the obtained standard deviation is about 3–5 times smaller. This can be explained by the fact that the printed images are more continuous, i.e., have a more uniformly distributed histogram, which makes the image synthesis more reliable. For Samsung, the dynamic range of the obtained deviation is about

³The code and configuration files are publicly available at https://gitlab. unige.ch/sip-group/stochastic-digital-twin



Fig. 6. Impact of the number of realizations on different metrics for the iPhone dataset.

 TABLE I

 Models performance on the iPhone data

TABLE II							
MODELS PERFORMANCE ON THE SAMSUNG DAT	ΓA						

Model	$\begin{array}{c} FID \\ \mathbf{x} {\rightarrow} \tilde{\mathbf{z}} \end{array}$	Hamm. dist.	$\begin{array}{c} FID\\ \mathbf{z}{\rightarrow}\tilde{\mathbf{x}} \end{array}$	MSE	SSIM	Model	$\begin{array}{c} FID\\ \mathbf{x}{\rightarrow}\tilde{\mathbf{z}} \end{array}$	Hamm. dist.	$\begin{array}{c} FID\\ \mathbf{z}{\rightarrow}\tilde{\mathbf{x}} \end{array}$	MSE	SSIM
W/O processing	289.68	0.300	289.68	0.254	0.249	W/O processing	381.44	0.314	381.44	0.278	0.193
pix2pix	11.82	0.232	11.64	0.005	0.910	pix2pix	8.53	0.241	20.18	0.004	0.908
CycleGAN	20.69	0.268	12.59	0.014	0.782	CycleGAN	8.85	0.283	22.85	0.015	0.694
TURBO _{CNN-RESNET-CNN}	6.56	0.239	10.20	0.005	0.915	TURBO ^{paired (w D)}	7.01	0.247	17.34	0.004	0.914
TURBO ^{paired (w/o D)}	35.35	0.210	12.22	0.004	0.925	TURBO	54.80	0.211	28.88	0.004	0.922
Palette _{mean}	4.64	0.211	9.00	0.004	0.915	Palette _{mean}	4.46	0.215	10.72	0.004	0.908

1.5 times smaller than for iPhone. For the scanner results one also observes a smaller amount of unreliable regions with the edge-transition regions being very well pronounced.

It was shown in [14] that the printed pixel's variability depends on the surrounding neighborhood that we refer to as *pattern* ω , where ω denotes a 3 \times 3 configuration of each pattern. To investigate if the same effect is present in the synthetically generated codes we define $2^{(3\times3)} = 512$ possible patterns and calculate the standard deviation of the central pixel for each of them through all generated codes. The results obtained for the iPhone dataset are shown in Fig. 4. For Samsung, we observed quite a similar picture. In Fig. 4 we can see the same tendency as in Fig. 3, namely, the general dynamic range of the standard deviation for the $\mathbf{x} \rightarrow \tilde{\mathbf{z}}$ channel is higher than for the $\mathbf{z} \to \tilde{\mathbf{x}}$, i.e., 0–0.45 versus 0.06–0.09. Also, we observe the pattern dependence for both channels but this dependence differs between the channels that is natural. In particular, in the $\mathbf{x} \to \tilde{\mathbf{z}}$ channel there are more patterns with the less variable central pixel, i.e., $\sigma(\omega)$ close to 0.

To study the similarity between the real CDP x and the synthetic counterparts \tilde{x} , we compute the probability of bitflipping for the central pixel for each pattern after Otsu binarization, as suggested in [14]. The results obtained for the iPhone are shown in Fig. 5⁴. We can see that some patterns almost certainly flip with $P_b(\omega)$ close to 1, whereas others

B. Aggregation techniques

For the Palette model, we study the impact of the number of realizations and different aggregation techniques. The iPhone results are shown in Fig. 6^5 , where Palette_{mean} denotes that the final output is obtained as a mean of predictions, while in Palette_{median} we take a median of predictions, and then the score is calculated for the aggregated prediction. In the case Palette_{mean of scores} the reference metric is calculated for each prediction and then the mean value of obtained scores is taken.

One can observe that $Palette_{mean of scores}$ error is almost constant, while the errors for $Palette_{mean}$ and $Palette_{median}$ are close to each other and decrease as the number of realizations increases. This can be explained by the fact that from one side the global error is the same at each realization but the local errors appear in different positions. The aggregation allows us to reduce them but, at the same time, it leads to the loss of stochasticity. In the same plot, one can see the results for pix2pix and TURBO^{paired (w/o D)} models, but these models are deterministic, and their results are not impacted by the number of realizations. Palette easily outperforms pix2pix on all metrics after 5-7 realizations. In terms of Hamming distance and SSIM Palette is not capable of outperforming

produce reliable results with $P_b(\omega)$ close to 0. But the most important thing is that the bit-flipping probability for the real x perfectly correlates with one for the synthetic \tilde{x} .

⁴For Samsung we observed the same tendency.

⁵The results for the Samsung data are similar to iPhone.

TABLE III MODELS PERFORMANCE ON THE SCANNER DATA

Model	FID x→ž	Hamm. dist.	$\begin{array}{c} FID\\ \mathbf{z}{\rightarrow}\tilde{\mathbf{x}} \end{array}$	MSE	SSIM
W/O processing	304.13	0.238	304.01	0.181	0.480
pix2pix	3.37	0.111	8.57	0.045	0.758
CycleGAN	3.87	0.155	4.45	0.049	0.732
TURBO ^{paired (w D) CNN-RESNET-CNN}	3.16	0.086	6.60	0.040	0.779
TURBO ^{paired (w/o D)}	6.21	0.100	28.110	0.036	0.778
Palette _{mean}	2.90	0.081	4.36	0.038	0.769

TURBO^{paired (w/o D)}. For MSE, Palette needs at least 10–20 realizations to surpass Turbo.

C. Models general performance

For further performance evaluation, we use Palette_{mean}. The inference time for a single realization of the Palette model applied to 280 test images is approximately 30 minutes. In contrast, the Turbo model achieves an inference time of just 15 seconds on the same GPU for the same dataset. Considering the inference time complexity, we found that using 21 realizations strikes a good balance between reasonable inference execution time, stochasticity preservation, and the accuracy of the final result. We compare the performance of the Palette in this configuration with the state-of-the-art pip2pix [8], Cycle-GAN [9], TURBO^{paired (w D)} and TURBO^{paired (w/o D)} on the same set of metrics as in [5]. W/O processing setup is used to estimate the baseline performance where we assume $\tilde{\mathbf{z}} = \mathbf{x}$ and $\tilde{\mathbf{x}} = \mathbf{z}$, i.e., an ideal printing-imaging channel without any distortions. The results obtained for the data enrolled by the iPhone and Samsung mobile phones are given in Tables I and II respectively.

It should be noted that the results obtained for the FID metric for the $\mathbf{x} \to \tilde{\mathbf{z}}$ and $\mathbf{z} \to \tilde{\mathbf{x}}$ channels are quite unstable and differ a lot between the models. This can be explained by the fact that FID was developed for natural images while, in our case, all CDP significantly differ from them. The other metrics demonstrate more coherent results.

For both mobile phones, the MSE results are almost identical for Palette and both Turbo configurations. In terms of SSIM, the results are also very close with a slight TURBO^{paired (W/o D)} superiority. TURBO^{paired (W/o D)} also outperforms the other models in terms of Hamming distance. The results for the remaining models exhibit marginal inferiority.

The result obtained for the data enrolled by the scanner are given in Table III and the general tendency is the same, namely, FID behavior is very unstable; CycleGAN demonstrates the worst results; the results for the Palette and Turbo are quite close but Palette is slightly superior to Turbo on the Hamming distance, albeit at the cost of significantly higher complexity of inference.

V. CONCLUSION

The current work is a continuation of our previous study [5] related to modeling of complex physical printing-imaging processes using a machine learning based models known as a *digital twin* for anti-counterfeiting applications based on CDP. The current work is dedicated to investigation of the applicability of DDPM for such modeling.

Our main interest was to explore the stochasticity of DDPM. The obtained results show that synthetic digital and CDP images are close enough to the real ones in terms of considered metrics. Moreover, the synthetic CDP images produced by DDPM fully reflect the natural randomness of the printing process. This makes the DDPM-based model a suitable candidate for the role of a synthetic generator. The general performance of the studied Palette model is comparable to that of the Turbo framework. The main drawback of the DDPM is the computation complexity of the inference stage.

The investigation of more advanced sampling techniques at the training and inference stages for the improvement of DDPM complexity is the main direction for our future work.

REFERENCES

- J. Picard, "Digital authentication with copy-detection patterns," in *Optical Security and Counterfeit Deterrence Techniques V*, vol. 5310. SPIE, 2004, pp. 176–183.
- [2] J. Picard, P. Landry, and M. Bolay, "Counterfeit detection with qr codes," in *Proceedings of the 21st ACM Symposium on Document Engineering*, 2021, pp. 1–4.
- [3] R. Yadav, I. Tkachenko, A. Trémeau, and T. Fournel, "Copy sensitive graphical code estimation: Physical vs numerical resolution," in 2019 IEEE International Workshop on Information Forensics and Security (WIFS). IEEE, 2019, pp. 1–6.
- [4] A. P. Ho, B. M. Hoang, W. Sawaya, and P. Bas, "Authentication using graphical codes: Optimisation of the print and scan channels," in 2014 22nd European Signal Processing Conference (EUSIPCO). IEEE, 2014, pp. 800–804.
- [5] Y. Belousov, B. Pulfer, R. Chaban, J. Tutt, O. Taran, T. Holotyak, and S. Voloshynovskiy, "Digital twins of physical printing-imaging channel," in *IEEE International Workshop on Information Forensics and Security* (*WIFS*), Shanghai, China, December 2022.
- [6] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," Advances in Neural Information Processing Systems, vol. 33, pp. 6840– 6851, 2020.
- [7] P. Dhariwal and A. Nichol, "Diffusion models beat gans on image synthesis," *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.
- [8] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," *CVPR*, 2017.
- [9] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proceedings* of the IEEE international conference on computer vision, 2017, pp. 2223–2232.
- [10] C. Luo, "Understanding diffusion models: A unified perspective," arXiv preprint arXiv:2208.11970, 2022.
- [11] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, "Palette: Image-to-image diffusion models," in ACM SIGGRAPH 2022 Conference Proceedings, 2022, pp. 1–10.
- [12] R. Chaban, O. Taran, J. Tutt, T. Holotyak, S. Bonev, and S. Voloshynovskiy, "Machine learning attack on copy detection patterns: are 1x1 patterns cloneable?" in 2021 IEEE International Workshop on Information Forensics and Security (WIFS). IEEE, 2021, pp. 1–6.
- [13] R. Chaban, O. Taran, J. Tutt, Y. Belousov, B. Pulfer, T. Holotyak, and S. Voloshynovskiy, "Printing variability of copy detection patterns," in *IEEE International Workshop on Information Forensics and Security* (WIFS), Shanghai, China, December 2022.
- [14] J. Tutt, O. Taran, R. Chaban, B. Pulfer, Y. Belousov, T. Holotyak, and S. Voloshynovskiy, "Mathematical model of printing-imaging channel for blind detection of fake copy detection patterns," in *IEEE International Workshop on Information Forensics and Security (WIFS)*, Shanghai, China, 12 2022.