

Assessing the Viability of Synthetic Physical Copy Detection Patterns on Different Imaging Systems

Roman Chaban, Brian Pulfer and Slava Voloshynovskiy
 Department of Computer Science, University of Geneva, Switzerland
 {roman.chaban, brian.pulfer, svolos}@unige.ch

Abstract—This paper explores the potential of synthetic physical Copy Detection Patterns (CDP) to improve the robustness of anti-counterfeiting systems. By leveraging synthetic physical CDP \hat{x} , we aim at enhancing security and cost-effectiveness across various real-world applications. Our research demonstrates that synthetic CDP offer substantial improvements in authentication accuracy compared to one based on traditional digital templates t . We conducted extensive tests using both a scanner and a diverse range of mobile phones, validating our approach through ROC analysis. The results indicate that synthetic CDP can reliably differentiate between original and fake samples, making this approach a viable solution for real-world applications, though requires an additional research to make this technology scalable across a variety of imaging devices.

Index Terms—Copy detection patterns, machine learning fakes, synthetic CDP, ROC analysis, mobile imaging.

I. INTRODUCTION

Copy Detection Patterns (CDP) have emerged as a promising anti-counterfeiting technology due to their cost-effectiveness, scalability, and ease of integration across various industries [1], [2]. These patterns are applicable to packaging, security labels, pharmaceuticals, currency, and identification documents, with their effectiveness stemming from their resistance to replication, machine verifiability, and versatility.

However, the sophisticated machine learning attacks have proven to be highly efficient against existing CDP, compromising their security [3]–[6]. These attacks utilize advanced algorithms to create high-quality fake CDP f that can bypass existing authentication methods, raising concerns about CDP’s non-cloneability, particularly in critical sectors like pharmaceuticals. Meanwhile, authentication based on reference CDP x_e , a.k.a. physical template, offers superior protection in contrast to one based on digital template t [7] but requires the enrollment of x_e , which compromises cost-effectiveness and should be used only in very specific applications.

To balance cost and robustness to cloneability (performance in the face of advanced fakes), this study endorses the use of synthetic physical CDP \hat{x} [8]–[10], which is obtained through a trained deep mapper. This approach mitigates the high acquisition costs of x_e while improving the low authentication accuracy associated with digital template t . Synthetic \hat{x} combines physical and digital benefits, providing a secure and cost-efficient solution. Building upon earlier works [7], [10], our research primarily evaluates the effectiveness of synthetic

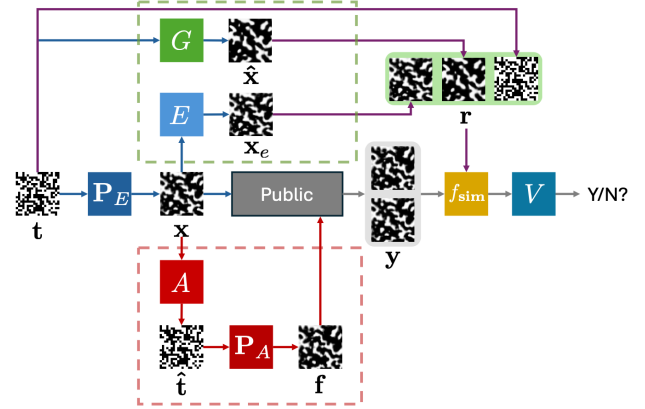


Fig. 1: The diagram demonstrating the real-world CDP lifecycle which includes the authentication of probe y , the possible attack scenario (red dashed line) yielding fake f , and the proposed defensive strategies (green dashed line) based on the reference $r \in \{t, x_e, \hat{x}\}$.

CDP \hat{x} against advanced machine-learning fakes f [3] through the scope of consumer mobile phones.

The primary focus is on a comparative analysis of the proposed authentication scenarios across different imaging systems. We conduct an analysis using both scanners and mobile phones to provide a comprehensive overview of various real-world conditions, utilizing a new dataset. Scanner, known for its high precision and consistency, provides reliable data for evaluation but is not considered as the practical use-case device for CDP verification. Thus, we test a range of mobile phones from older iPhone XS to cutting-edge iPhone 15 Pro Max. This diversity allows us to assess the authentication performance across multiple imaging systems, which vary significantly in terms of optical capabilities, sensor size, focal length, field of view, etc.. By including a broad spectrum of devices in our study, we aim to demonstrate that the authentication based on \hat{x} can perform reliably under various conditions, while at the same given conditions the authentication based on t fails.

In our study, we pursue the following goals:

- 1) to conduct the experiments using two industrial digital offset printers, which are a promising printing technology for CDP in view of the increasing demand for variable printing and product personalization;
- 2) to underline the vulnerabilities and limitations of nowadays authentication methodologies based on digital templates t ;

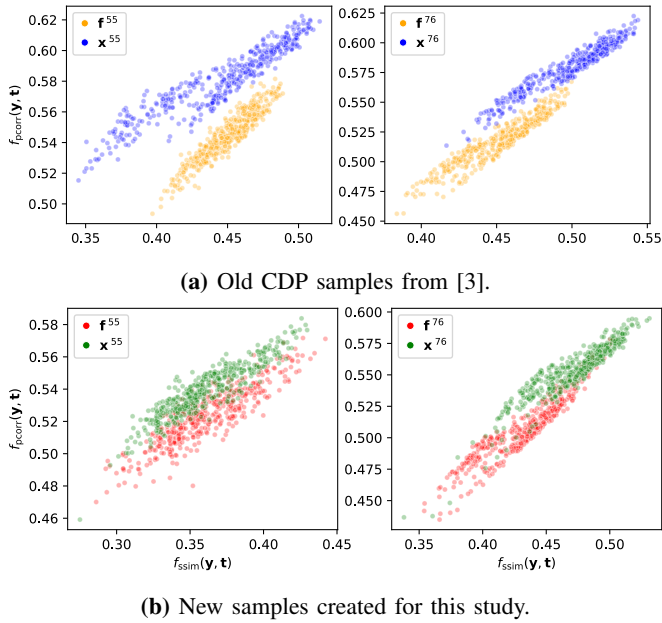


Fig. 2: The distributions of $f_{\text{Pcorr}}(\mathbf{y}, \mathbf{t})$ and $f_{\text{Ssim}}(\mathbf{y}, \mathbf{t})$ for both printers, where probes \mathbf{y} were scanned with a flatbed scanner. \mathbf{f} and \mathbf{x} denote fake and original probes, respectively, and \mathbf{t} stands for the digital template.

- 3) to demonstrate the impact of different mobile phones on the authentication performance and on par with the flatbed scanner;
- 4) to compare the existing authentication methods based either on digital template \mathbf{t} or on physical template \mathbf{x}_e to the proposed one based on synthetic physical CDP $\hat{\mathbf{x}}$;
- 5) to outline the applicability scope of authentication based on $\hat{\mathbf{x}}$.

Meanwhile, we are aware that in such a study there is a multitude of possibilities and subjects for investigation, thus we have imposed the following limitations to deliver a clear message and for the sake of simplicity:

- 1) We conduct experiments in an imaging-consistent manner, meaning that verification probe \mathbf{y} is compared to reference $\mathbf{r} \in \{\mathbf{x}_e, \hat{\mathbf{x}}\}$ of the same imaging system.
- 2) We consider the similarity assessment via two similarity metrics f_{sim} : Pearson Correlation Coefficient (f_{Pcorr}) and Structural Similarity Measure Index [11] (f_{Ssim}), since the cited studies have proven their effectiveness and to facilitate fair comparison and reproducibility.
- 3) We choose iPhones since they provide reliable and consistent imaging across their device lineup, enabling a better track of imaging evolution.

In conclusion, synthetic physical CDP $\hat{\mathbf{x}}$ offers a promising solution to the vulnerabilities in traditional CDP systems by leveraging both physical and digital advantages for enhanced security and cost-effectiveness. Our study, involving scanner and mobile phone evaluations, aims to validate this solution and provide a robust framework for its implementation across various industries. The findings will contribute to the

development of more secure anti-counterfeiting technologies, protecting critical products and documents.

II. DATASET

In this study, we introduce a novel dataset created using previously established methodologies and CDP design principles. The dataset was produced on two industrial digital offset printers: HP Indigo 5500 (HPI55) and 7600 (HPI76). Unlike previous studies [3], [7], [10], the resulting difference between original and fake CDP is even less distinguishable both perceptually and statistically as demonstrated in Fig. 2. This in turn presents a more challenging problem for the classification¹.

The dataset consists of 144 unique CDP, each randomly generated with a binary distribution where 50% of the pixels are black and 50% are white. Such a decision was based on prior research indicating that this amount of CDP suffices to assess the general statistics of CDP. Consequently, the study involves 576 physical instances, derived from 2 printers \times 2 origins (originals and fakes) \times 144 unique CDP.

A key highlight of our new dataset collection is the number of used mobile phones, from the nearly obsolete iPhone XS (*XS wide*) to the most advanced models currently available, such as the iPhone 15 Pro Max (*15 wide, 15 macro*), which also includes the iPhone 12 Pro (*12 wide*) and iPhone 14 Pro (*14 wide, 14 macro*). Notably, the iPhone 14 Pro and iPhone 15 Pro Max feature two relevant camera modules for this study: a regular wide and an ultra-wide macro lens. The macro lens, despite being a consumer-level component, already yields a much higher effective resolution, hypothetically surpassing the quality of Epson V850 (*Epson*), the scanner used in our study. As shown in Fig. 4 the samples acquired with different imaging systems clearly possess distinguishable resulting statistics of similarity metric $f_{\text{sim}}(\mathbf{y}, \mathbf{t})$, however, we can observe a certain overlap between *14 wide* and *15 wide* systems, and for *14 macro* and *15 macro*.

Each physical CDP, whether fake or original, was captured 3 times by every imaging device. After automatic quality control, a fraction of $\approx 5\%$ photos was discarded. This approach ensures high-quality images by excluding blurred or low-contrast captures. The effective PPI and a scale factor with respect to \mathbf{t} differ between imaging systems. For authentication purposes, originals were divided into enrollment \mathbf{x}_e and verification probes \mathbf{y} .

The dataset was acquired using a custom-built mobile application, which enabled the capture of uncompressed images. This application was designed to obtain photos of the highest possible quality by:

- 1) mitigating most causes of variability, such as illumination changes, motion blur, and out-of-focus issues;
- 2) capturing from the minimal possible focal distance to yield the highest effective PPI;
- 3) minimization of perspective transformation by guidance in the user interface.

¹Dataset is available here: <https://github.com/romaroman/cdp-synthetics-dataset>.




Fig. 3: The examples of the left-upper crop of CDP #166 printed on HPI55 and captured with a whole range of studied devices. Along with the digital template \mathbf{t} and its printed form \mathbf{x} , the figure comprises the synthetic physical $\hat{\mathbf{x}}$ that visually closely resembles real physical CDP \mathbf{x} .




Fig. 4: The scatter plots of both f_{sim} for all given combinations of printer and imaging device.

All obtained CDP were aligned with respect to the digital template using SIFT descriptors [12], followed by pixel-precision alignment using the same digital template.

In summary, our dataset, printed using HPI55 and HPI76 printers, and captured with six different mobile imaging systems, presents a more challenging authentication dataset by reducing the statistical distinguishability between original and fake CDP. This dataset, acquired by a number of modern consumer-grade mobile phones and a specialized application, will serve as a valuable resource for further studies on CDP-based authentication systems.

III. METHODOLOGY

In this section, we outline the methodology used in this study for the CDP authentication. The described steps are independently applied to each unique combination of industrial printer and imaging device.²

A. General Overview

To create the original CDP \mathbf{x} , we begin by printing the digital template \mathbf{t} . To generate the fake CDP \mathbf{f} , we use the original CDP \mathbf{x} as input to a pre-trained U-Net [3], [13]

estimator A , which produces an estimated digital template $\hat{\mathbf{t}}$ as shown by the red dashed line in Fig. 1. This estimated template $\hat{\mathbf{t}}$ is then reprinted using the same printer and settings, resulting in the fake CDP \mathbf{f} . The complete description of fakes \mathbf{f} production details can be found in [3].

The original CDP \mathbf{x} is divided into two subsets. The first subset is used both for enrollment denoted as \mathbf{x}_e and for training the pix2pix model [14] to generate synthetic physical CDP $\hat{\mathbf{x}}$. Effectively, for the training of synthetic generator G only approximately 30 CDPs are sufficient, and such generator is well-scalable for much bigger volumes [9].

The second subset of \mathbf{x} is reserved for use as incoming probes \mathbf{y} during authentication. These probes \mathbf{y} (which can be either original \mathbf{x} or fake \mathbf{f}) are compared against references \mathbf{r} , (which can be digital template \mathbf{t} , synthetic physical $\hat{\mathbf{x}}$, or enrolled \mathbf{x}_e CDPs) to calculate the similarity score $f_{\text{sim}}(\mathbf{y}, \mathbf{r})$.

B. Training of Synthetic CDP Generator

The pix2pix model [14] G is used to learn the mapping from the digital template CDP \mathbf{t} , to the synthetic physical CDP $\hat{\mathbf{x}}$. We employed a state-of-the-art paired pix2pix architecture, where a U-Net [13] serves as the generator G and PatchGAN is used as the discriminator D . The training process involves optimizing specific loss functions for both the discriminator

²GitHub repository: <https://github.com/romaroman/cdp-synthetics>.




Fig. 5: The distributions of $f_{\text{pcorr}}(\mathbf{y}, \mathbf{r})$ for CDP captured with the worst-case imaging system, *iPhone XS wide*.

and the generator to achieve high-quality synthetic CDP for pairs $\{\mathbf{t}_i, \mathbf{x}_i\}_{i=1}^N$, where $N = 40$ is the number of training pairs.

The discriminator D is trained to distinguish between real physical CDP \mathbf{x}_e and synthetic physical CDP $\hat{\mathbf{x}}$ generated by the generator G . Its loss function is defined as:

$$\mathcal{L}_D = -\mathbb{E}_{(\mathbf{t}, \mathbf{x})} [\log D(\mathbf{t}, \mathbf{x})] - \mathbb{E}_{\mathbf{t}, \hat{\mathbf{x}}=G(\mathbf{t})} [\log(1 - D(\mathbf{t}, \hat{\mathbf{x}}))], \quad (1)$$

where $\mathbb{E}_{(\mathbf{t}, \mathbf{x})} [\bullet]$ denotes the mathematical expectation.

This loss function encourages D to assign high probabilities to real CDP and low probabilities to synthetic ones, thereby becoming sensitive to identifying real CDP.

Conversely, G is trained to produce synthetic physical $\hat{\mathbf{x}}$ that are indistinguishable from \mathbf{x} . Its loss function is given by:

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{t}, \hat{\mathbf{x}}=G(\mathbf{t})} [\log D(\mathbf{t}, \hat{\mathbf{x}})] + \lambda \mathbb{E}_{(\mathbf{t}, \mathbf{x})} [\|\mathbf{x} - G(\mathbf{t})\|_1], \quad (2)$$

where λ controls the trade-off between two terms.

This loss function encourages G to generate samples that D misclassifies as real, while also maintaining a high degree of visual and structural fidelity to \mathbf{x} by including a similarity term weighted by λ .

By balancing these loss functions, the pix2pix model effectively learns to generate $\hat{\mathbf{x}}$ that are both statistically and visually similar to \mathbf{x} . The generator's output is thus guided to not only deceive the discriminator but also to closely resemble real CDP, ensuring that the synthetic CDP can be used reliably in place of original CDP for authentication purposes.

For model training we have conducted a train/test split by ensuring that the model trains on a small fraction ($\approx 30\%$) of \mathbf{x}_e and is tested on the previously unseen \mathbf{t} , effectively generating $\hat{\mathbf{x}}$ which were never subjected to loss computation and model optimization.

This approach ensures that the pix2pix model effectively learns to generate high-quality synthetic CDP that are not only visually indistinguishable from real CDP but also maintain a high degree of similarity from a statistical point of view.

C. Similarity Assessment

In the authentication stage, we compute the similarity metrics f_{pcorr} and f_{ssim} , for the incoming probe $\mathbf{y} \in \{\mathbf{x}, \mathbf{f}\}$ with respect to the reference $\mathbf{r} \in \{\mathbf{t}, \hat{\mathbf{x}}, \mathbf{x}_e\}$. These similarity scores are then analyzed using Receiver Operating Characteristic (ROC) curves, and the corresponding Area Under the Curve (AUC) values are calculated to evaluate the effectiveness of each reference \mathbf{r} . We do not create a specific threshold-based classifier, as the ROC analysis provides a sufficient measure of potential classification performance in our study. Obviously, in practical use cases, an appropriate decision threshold should be chosen based on the provided ROC curves and the desired probability of false acceptance and probability of miss.

IV. RESULTS

Referring back to Fig. 4, we conducted same-device experiments, where the incoming probe \mathbf{y} is compared to \mathbf{x}_e or $\hat{\mathbf{x}}$ from the same physical and acquisitional origin. Future work should explore cross-device testing, as some imaging systems (like *14 macro* and *15 macro*) that share similar statistical distributions as claimed by the vendor.

As expected from the distributions shown in Fig. 2, we observed similar degrees of separability in terms of $f_{\text{sim}}(\mathbf{y}, \mathbf{t})$ for mobile phones. This is well-supported in Fig. 5 (Digital \mathbf{t} plots), and the trend is consistent for both printers. The AUC values indicated in the plots' titles correspond to the ROC curves in Fig. 6. Table I presents a comprehensive comparison of $f_{\text{sim}}(\mathbf{y}, \mathbf{t})$ performance across all imaging systems. Key conclusions for $f_{\text{sim}}(\mathbf{y}, \mathbf{t})$ are:

- 1) Authentication accuracy increases with imaging effective resolution (from old devices to more recent ones).

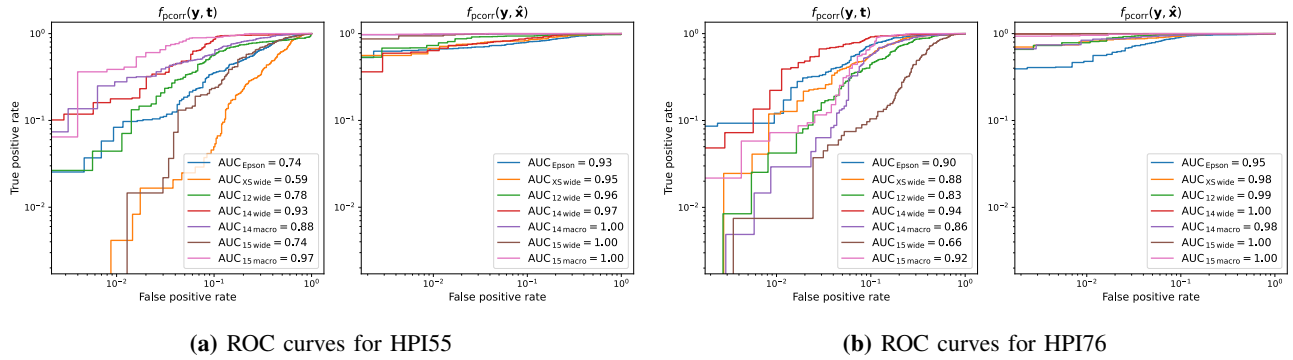


Fig. 6: ROC curves for $f_{\text{pcorr}}(\mathbf{y}, \mathbf{r})$, where $\mathbf{r} \in \{\mathbf{t}, \hat{\mathbf{x}}\}$, with corresponding AUC values for each imaging system. \mathbf{x}_e is not shown as its AUC equals 1.00 for all systems. $\hat{\mathbf{x}}$ were trained for each imaging system.

\mathbf{r}	f_{sim}	HPI55							HPI76						
		<i>Epson</i>	<i>XS wide</i>	<i>12 wide</i>	<i>14 wide</i>	<i>15 wide</i>	<i>14 macro</i>	<i>15 macro</i>	<i>Epson</i>	<i>XS wide</i>	<i>12 wide</i>	<i>14 wide</i>	<i>15 wide</i>	<i>14 macro</i>	<i>15 macro</i>
\mathbf{t}	f_{pcorr}	0.74	0.59	0.78	0.93	0.74	0.88	0.97	0.90	0.88	0.83	0.94	0.66	0.86	0.92
	f_{ssim}	0.51	0.71	0.68	0.90	0.64	0.74	0.79	0.81	0.78	0.69	0.93	0.63	0.81	0.81
$\hat{\mathbf{x}}$	f_{pcorr}	0.93	0.95	0.96	0.97	1.00	1.00	1.00	0.95	0.98	0.99	1.00	1.00	0.98	1.00
	f_{ssim}	0.93	0.96	0.97	0.96	1.00	0.99	1.00	0.94	0.98	0.99	1.00	1.00	0.97	1.00
\mathbf{x}_e	f_{pcorr}	1.00	1.00	0.97	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	f_{ssim}	1.00	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

TABLE I: The breakdown of AUC values presented for the whole range of similarity metrics, references, imaging systems and printers of current study scope.

- 2) Old phones demonstrate unsatisfactory performance for the authentication based on digital template \mathbf{t} .
- 3) Phones equipped with macro lenses sometimes outperform scanners, approaching ideal performance.
- 4) Mobile phone sampling shows high variability and uncertainty, causing considerable fluctuations in AUC.
- 5) On average, f_{ssim} is less discriminative than f_{pcorr} .
- 6) Neither printer nor imaging system combination provides satisfactory AUC.

Next, we assessed the performance of $f_{\text{sim}}(\mathbf{y}, \mathbf{x}_e)$ as the physical baseline reference. As shown in Fig. 5 (Physical \mathbf{x}_e plots), even the oldest mobile phone shows negligible overlap between \mathbf{x} and \mathbf{f} for both printers, underscoring the superiority of $f_{\text{sim}}(\mathbf{y}, \mathbf{x}_e)$. This trend is consistent across all acquisitions. When $\mathbf{y} = \mathbf{x}$, we effectively compute $f_{\text{sim}}(\mathbf{x}, \mathbf{x}_e)$, where both \mathbf{x} and \mathbf{x}_e originate from the same physical CDP, achieving the highest mutual information. Consequently, $f_{\text{sim}}(\mathbf{f}, \mathbf{x}_e)$ cannot match this level of microstructural similarity. Table I shows perfect separability for any printer and imaging device pair, except for *12 wide* due to alignment issues discovered through manual investigation. Therefore, ROC curves for these results are redundant and not shown in this paper. Key findings can be summarized as:

- 1) Results align well with previous studies [7] and are explained by \mathbf{x}_e being a physically unclonable function [15] of CDP.
- 2) $f_{\text{sim}}(\mathbf{y}, \mathbf{x}_e)$ represents a sufficient statistics for authenti-

cation systems.

- 3) Authentication based on physical CDP \mathbf{x}_e achieves excellent results even for *XS wide*. However, such an approach requires enrollment from each physical CDP \mathbf{x} .

With these baselines, we assess the viability of synthetic physical $\hat{\mathbf{x}}$. No enrollment from all future physical objects is necessary, once the generator G is trained on a small fraction of enrolled physical \mathbf{x}_e . Fig. 5 (Synthetic $\hat{\mathbf{x}}$ plots) shows significant improvement in AUC for $f_{\text{pcorr}}(\mathbf{y}, \hat{\mathbf{x}})$ compared to $f_{\text{pcorr}}(\mathbf{y}, \mathbf{t})$, regardless of the printer. This trend is consistent across all imaging devices, as shown in Table I and Fig. 6. In many scenarios, $\hat{\mathbf{x}}$ achieves perfect separability (AUC = 1), though this might be due to a small sample size and may decrease with a larger number of samples. Key takeaways are:

- 1) $\hat{\mathbf{x}}$ as reference \mathbf{r} significantly increases separability of \mathbf{x} from \mathbf{f} for low quality imaging devices.
- 2) For scenarios where AUC < 1, improvements such as fine-tuning, data refinement, and tweaking network architecture or loss functions are still possible.
- 3) Although $\hat{\mathbf{x}}$ lags behind \mathbf{x}_e in terms of AUC, it is sufficient for certain imaging systems in its current implementation.
- 4) AUC for scanner-based results is the lowest compared to any mobile phone, suggesting that the synthetic generator G is more suited for the mobile phone images.
- 5) Slightly better performance of f_{ssim} for HPI55 with *12 wide* pair may be linked to the generator G 's loss

function, which includes SSIM, L2, and L1 norms.

In summary, the results highlight the separability of original \mathbf{x} and fake \mathbf{f} samples across various imaging and printing systems. $f_{\text{pcorr}}(\mathbf{y}, \hat{\mathbf{x}})$ yields higher AUC values compared to $f_{\text{sim}}(\mathbf{y}, \mathbf{t})$. Mobile phones show high variability in AUC, with macro lenses performing close to ideal in some cases. The $f_{\text{sim}}(\mathbf{y}, \mathbf{x}_e)$ approach consistently achieves perfect separability due to the physically unclonable nature of \mathbf{x}_e . The synthetic reference $\hat{\mathbf{x}}$ significantly enhances separability, though still below \mathbf{x}_e 's performance. Scanner-based results exhibit the lowest AUC, indicating that the generator G adapts better to mobile phone conditions. Overall, these findings underscore the potential for improved authentication using \mathbf{x}_e and $\hat{\mathbf{x}}$ references, with implications for low-end device applications. The obtained results are based on the pix2pix model, which is a component of the more powerful TURBO model [16], known for its superior performance in CDP applications [9]. Therefore, we anticipate even better results with an increase in model complexity.

V. CONCLUSIONS AND FUTURE WORK

In this study, we have explored the potential of synthetic physical $\hat{\mathbf{x}}$ as a viable alternative to the existing authentication framework. Our research demonstrates that synthetic physical $\hat{\mathbf{x}}$ significantly improves authentication accuracy compared to digital templates \mathbf{t} , with consistent performance across various imaging systems.

Our findings indicate that synthetic physical $\hat{\mathbf{x}}$ can substantially enhance the separability between original \mathbf{x} and fake \mathbf{f} samples. While mobile phones exhibit some variability in authentication performance, the use of ultra-wide macro lenses pushes the frontier of mobile phone-based authentication of CDP by showing near-ideal results in certain cases. Although $\hat{\mathbf{x}}$ does not yet match the perfect performance of physical enrollment \mathbf{x}_e , it demonstrates significant potential. Further refinements in model training, data augmentation, and network architecture could enhance its performance even further.

The study supports the adoption of synthetic $\hat{\mathbf{x}}$ over digital templates \mathbf{t} or physical enrollment \mathbf{x}_e in scenarios where cost-effectiveness and scalability are critical. This research validates that $\hat{\mathbf{x}}$ can perform reliably under various conditions, offering a robust framework for real-world anti-counterfeiting applications across different industries.

Future work will include comprehensive cross-device testing to investigate the robustness of synthetic $\hat{\mathbf{x}}$. This will involve examining how estimators trained on one mobile phone model perform with probes from different systems. We will also explore the integration of other similarity metrics and advanced machine-learning techniques to enhance the accuracy and reliability of CDP-based authentication. Expanding the dataset and incorporating more diverse imaging systems will provide further insights into the practical applications of synthetic CDP.

Overall, this research contributes to the development of more secure and economical anti-counterfeiting technologies, providing a valuable resource for the protection of critical

products. The use of old mobile phones might be of interest to developing countries, where the fraction of fakes is relatively high and the mobile phone market is dominated by older models.

REFERENCES

- [1] J. Picard, "Digital authentication with copy-detection patterns," *Electron. Imaging*, vol. 5310, 06 2004.
- [2] —, "Copy detectable images: from theory to practice," *NIP24: international conference on digital printing technologies and digital fabrication*, pp. 796–798, 2008.
- [3] R. Chaban, O. Taran, J. Tutt, T. Holotyak, S. Bonev, and S. Voloshynovskiy, "Machine learning attack on copy detection patterns: are 1×1 patterns cloneable?" in *2021 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2021, pp. 1–6.
- [4] O. Taran, S. Bonev, and S. Voloshynovskiy, "Clonability of anti-counterfeiting printable graphical codes: A machine learning approach," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 2482–2486.
- [5] R. Yadav, I. Tkachenko, A. Trémeau, and T. Fournel, "Estimation of copy-sensitive codes using a neural approach," in *Proceedings of the ACM Workshop on Information Hiding and Multimedia Security*, ser. IH&MMSec'19. New York, NY, USA: Association for Computing Machinery, 2019, p. 77–82. [Online]. Available: <https://doi.org/10.1145/3335203.3335718>
- [6] E. Khermaza, I. Tkachenko, and J. Picard, "Can copy detection patterns be copied? evaluating the performance of attacks and highlighting the role of the detector," in *2021 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2021, pp. 1–6.
- [7] R. Chaban, E. Ictet, O. Taran, and S. Voloshynovskiy, "Comparative analysis of copy detection patterns and physical unclonable functions in authentication systems: A mobile phone perspective," *2024 32nd European Signal Processing Conference (EUSIPCO)*, August 2024.
- [8] Y. Belousov, B. Pulfer, R. Chaban, J. Tutt, O. Taran, T. Holotyak, and S. Voloshynovskiy, "Digital twins of physical printing-imaging channel," in *2022 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2022, pp. 1–6.
- [9] Y. Belousov, G. Quéant, B. Pulfer, R. Chaban, J. Tutt, O. Taran, T. Holotyak, and S. Voloshynovskiy, "A machine learning-based digital twin for anti-counterfeiting applications with copy detection patterns," *IEEE Transactions on Information Forensics and Security*, vol. 19, pp. 3395–3408, 2024.
- [10] B. Pulfer, Y. Belousov, J. Tutt, R. Chaban, O. Taran, T. Holotyak, and S. Voloshynovskiy, "Anomaly localization for copy detection patterns through print estimations," in *2022 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2022, pp. 1–6.
- [11] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [12] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [13] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [14] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1125–1134.
- [15] R. Maes, *Physically Unclonable Functions: Constructions, Properties and Applications*. Springer Berlin Heidelberg, 2013.
- [16] G. Quéant, Y. Belousov, V. Kinakh, and S. Voloshynovskiy, "Turbo: The swiss knife of auto-encoders," *Entropy*, vol. 25, no. 10, 2023. [Online]. Available: <https://www.mdpi.com/1099-4300/25/10/1471>