Embedding Robust Watermarking into Pattern to Protect the Copyright of Ceramic Artifacts

Lei Tan[1], Yuliang Xue[1], Guobiao Li[1], Zhenxing Qian[1]*, Sheng Li[1], Chunlei Bao[2]

[1]School of Computer Science, Fudan University
[2]The Arts and Technology Education Centre, Fudan University
leitan23@m.fudan.edu.cn, zxqian@fudan.edu.cn

Abstract

Ceramic artworks with elegant patterns present enormous collectible value and profits. To claim the copyright, the builder usually pastes their conspicuous stamp on the bottom or side of the ceramic artworks, which inevitably affects the external image of the artwork. In addition, the stamp is weak in resisting forgery attacks due to its visible nature. To address the above issues, we propose in this paper a novel framework for embedding invisible watermarking into patterns of the ceramic artworks. In the framework, a template-based watermarking embedding scheme is designed to map the watermark to an invisible template, which is added to the ceramic pattern to create its watermarked version. A distortion layer is further proposed to model the distortion of ceramic patterns in the ceramic manufacturing process, where a color-halftoning and an adaptive brightness adjustment strategy are developed to counter the print and firing operations that introduce the most significant distortions. Finally, a deep decoder is learned to extract the watermarking from the distorted pattern. Various experiments have been conducted to demonstrate the advantage of our proposed method for protecting the copyright of the ceramic artworks, which provides reliable watermark extraction accuracy without the need for a conspicuous stamp.

Introduction

Ceramic artworks are made through a complex process using ceramic materials, including clay, inks, coatings, etc. Due to its elegant appearance, ceramic artwork is liked by many people around the world and is used to decorate their homes and cultivate their sentiments. Crafted ceramic artwork with exquisite patterns is extremely valuable and often becomes more precious over time. For instance, at a Christie's auction in London in 2005, a Chinese Yuan Dynasty blue-and-white vase decorated with the "Guigu's Descent from the Mountain" motif sold for £15.688 million. Tempted by the enormous profits, many pirates forge ceramic artworks that look like real ones. Between 2020 and 2023, over 1,500 registered cases of ceramic art forgery were reported in Jingdezhen, Jiangxi Province, China.

To claim the copyright, the builder usually pastes their conspicuous stamp on the bottom or side of the ceramic artwork, as shown in Fig. 1(a). However, the stamp may affect the external appearance of the artworks and the artistic conception expressed by their pattern. On the other hand, the stamps are also prone to be forged due to their visible nature. As a result, pirated ceramic artworks with fake stamps are rampant in the market.

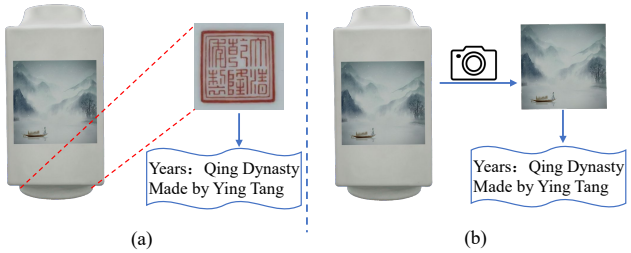

Figure 1: The traditional (a) and proposed (b) ways to protect the copyright of ceramic artworks.

Robust watermarking technology (Wang et al. 2024) provides a potential way to address the above limitations. It was used to protect the copyright of digital media by embedding watermarks within them. Inspired by this, we consider to embed the copyright data within the ceramic artwork's pattern in an imperceptible and robust way. 'imperceptible' refer to that the watermarked pattern looks similar with its clean version. 'robust' represent that the hidden data could be reliably extracted from the distorted pattern (i.e., the watermarked pattern suffers serious distortion during the ceramic manufacturing process). In this way, the created ceramic artworks with watermarked patterns not only retain their elegant image but also allow for the retrieval of hidden copyright data from their patterns when necessary, as shown in Fig. 1(b).

Unfortunately, existing image watermarking methods (Zhu et al. 2018; Tancik, Mildenhall, and Ng 2020) are designed either for digital or physical printed/screened and camera-captured images, which may be inadequate to provide reliable watermark retrieval accuracy in the ceramic artwork protection scenario. Specifically, digital image watermarking involves minor alteration to the original image and focuses on resisting distortions caused by digital image processing, such as Gaussian noise, Gaussian blur, JPEG

*Corresponding author.

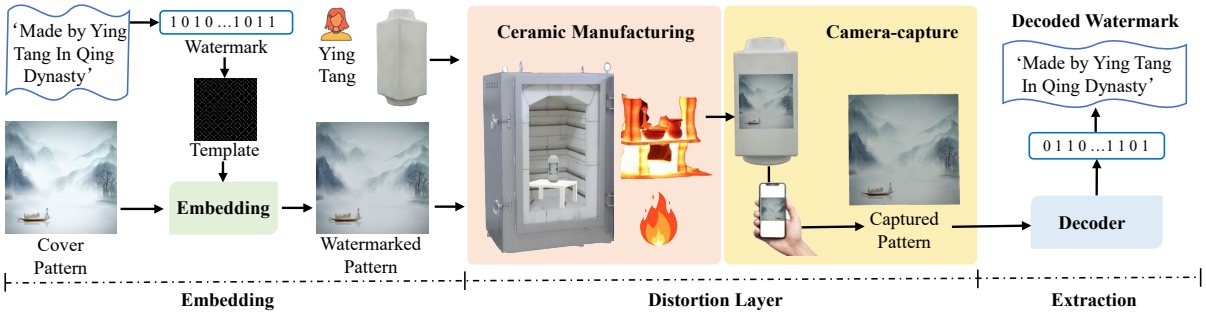

Figure 2: Illustration of the proposed framework.

compression, etc. It is vulnerable to severe physical distortion. Recent advanced physical image watermarking methods (Fang et al. 2018; Jia et al. 2022) that utilize deep neural networks (DNNs) demonstrate resistance to print-shooting or screen-shooting physical distortions. Most of them follow the encoder-distortion-layer-decoder architecture. Through deliberately designing the differentiable distortion layers to approach the space of distortions resulting from physical scenario, they jointly train the deep encoder and decoder to overfit the target physical distortion. As such, the embedded data by the encoder survives (i.e., be extracted by the decoder) in physical environments. However, the specificity of the distortion layer and the overfitting of the encoder-decoder weaken their generalizability. They usually fail to resist unknown distortions, e.g, the distortions in the ceramic manufacturing process.

Given that, we propose in this paper a novel framework to embed watermarking into patterns of the ceramic artworks, as shown in Fig.2. Unlike previous physical image watermarking framework that trains the deep encoder for data embedding, we design a template-based embedding module to map the watermark to an invisible template, which is added to the ceramic pattern to create its watermarked version. The strength of the template is adjustable to balance the visual quality of the watermarked pattern and the recovery accuracy of the hidden data. To accurate model the distortion in ceramic manufacturing environment, we deconstruct its process and analysis each steps. Based on the analysis, we develop a color-halftoning and an adaptive brightness adjustment strategy to counter the print and firing operations that introduce the most significant distortion. To make the watermark be extracted in camera-captured scenario, we further incorporate the transformations in real photography into our distortion layer. Finally, we train a deep decoder to extract the watermarking from the distorted ceramic patterns. Comprehensive experiments indicate the advantage of our method for protecting the copyright of ceramic artworks. It provides exceptional robustness, with approximately a 10% improvement in bit accuracy compared to the state-of-the-art (SOTA) physical image watermarking methods.

The main contributions are summarized below.

- We investigate the possibility that protecting the copyright of valuable ceramic artworks by embedding invisible watermarking into their patterns.
- We propose a novel robust watermarking framework in ceramic pattern decal-shooting scenario, which consists of a handcrafted template-based watermark embedding module, a tailored simulated distortion layer, and a learnable deep watermark extraction network.
- We design the distortion layer to approximate the space of distortions resulting from ceramic manufacturing environment.

Related Works

Digital Image Watermarking

Digital image watermarking aims to protect the copyright of digital image and are capable of resisting distortions caused by digital image processing, such as Gaussian noise, Gaussian blur, JPEG compression, etc. It typically embed watermarks in the spatial or frequency domain of the image. Akhaee et al. (Akhaee, Sahraeian, and Jin 2011) use the low-frequency components of image blocks for data hiding to obtain high robustness against attacks. Zareian et al. (Zareian and Tohidypour 2014) propose a novel quantization-based information hiding approach that is invariant to gain attacks. Later, with the rise of deep learning in computer vision, Hidden (Zhu et al. 2018) proposes the first end-to-end trainable framework, introducing a novel architecture that jointly optimizes the hiding and extraction components, enabling robust and intelligent information hiding. However, the aforementioned methods make minor modifications to digital images and are susceptible to significant physical distortions.

Physical Image Watermarking

Physical image watermarking is more robust than digital image watermarking, and the embedded data by them could survive in screen-shooting or print-shooting environment. Earlier studies designs hand-crafted algorithm to hide robust watermark within images. For instance, the works (Zheng, Zhao, and El Saddik 2003; Lin et al. 2001) embed watermarks into the Fourier-Mellin transform domain of images to resist the rotation, scaling, and translation distortions. Fang et al. (Fang et al. 2018) propose an intensity-based scale-invariant feature transform algorithm to resist screen-shooting distortion. Recent advanced physical image watermarking methods take advantage of DNNs for data embedding and extraction in an end-to-end manner, which achieve

superior performance than hand-crafted ones. They typically follow the encoder-distortion-layer-decoder architecture.

StegaStamp (Tancik, Mildenhall, and Ng 2020) develops a distortion layer that learns a robust encoding and decoding algorithm, which aims to enhance the robustness against print-shooting perturbations. Different from StegaStamp, LFM (Wengrowski and Dana 2019) trains a neural network to learn the distortion of camera-display transfer as the distortion layer. PIMoG (Fang et al. 2022) summarizes the most influenced distortions of the screen-shooting process, which performs effective results in different screen-shooting conditions. DeNoL (Fang et al. 2023) further simulates a decoupling distortion layer for cross-channel simulation which only needs few-shot samples. LIM (Jia et al. 2022) learns invisible markers and developed a localization module, which considerably reduces the time of correcting geometric distortions without breaking the visual invisibility. To adapt the film-coating scenario, WRAP (Liu et al. 2023) introduces the film-coating simulation network as a crucial component of the distortion layer, improving the watermarking embedding and extraction network robust against film-coating applied to printed photographs. The design of the (differentiable) distortion layer is crucial to the aforementioned method. Generally, the closer the simulated distortion layer is to the distortion space caused by real scenario, the more robust the trained encoder and decoder will be. However, despite being highly robust to the simulated distortions used in training, their performance often falls short when dealing with unknown physical distortions.

In this paper, we explore the feasibility that embedding watermarking into ceramic artworks' patterns to protect their copyright. To achieve this, we design a novel framework for the ceramic pattern decal-shooting scenario which has not been considered. Drawing on the successful experiences of the advanced physical image watermarking, we carefully design a specialized distortion layer to approach the distortions resulting from the ceramic manufacturing process.

Preliminary

Ceramic Manufacturing Process

The ceramic manufacturing process can be broadly divided into two categories: hand-painting and decal. The former refers to painting patterns on the fired ceramic with artists' hands, while the later is to pasting the printed pattern onto ceramics and then firing it. In this paper, we focus on the latter and detail its steps below. The illustration is shown in Fig. 3.

- Printing. The printing stage refers to printing a digital image onto glaze transfer paper using a printer that contains glaze ink.
- Spraying Cover-oil. It indicates spraying cove-roil onto the surface of the printed paper to fuse colors onto the glaze, so the ceramic decoration becomes durable.
- Pasting. The glaze transfer paper is soaked in water to separate the cover-oil with ink from the paper, which is then pasted to the surface of the ceramic.

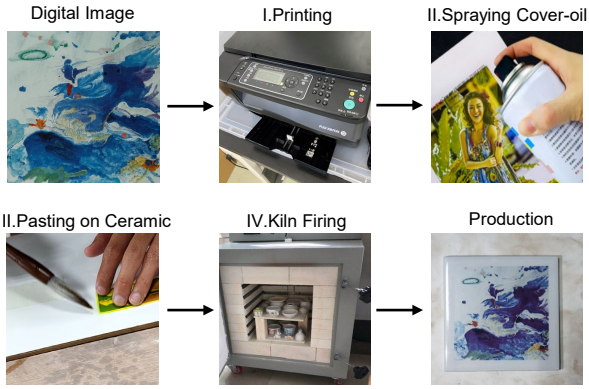

Figure 3: Illustration of ceramic manufacturing process.

- Firing. This stage aims to fire the ceramic body pasted with glaze paper at a higher temperature, which fuses the vivid pattern onto the ceramic surface.

The ceramic manufacturing process introduces numerous perturbations to the image, affecting the robustness of the watermarking schemes, where the most important effects are printing and firing. In this paper, we develop a color-halftoning and an adaptive brightness adjustment to simulate the distortion of the ceramic manufacturing process, respectively.

Camera-capture Process

The camera-capture process is briefly summarized in three steps, including light striking the image sensor, converting the light signals into electrical signals and then into digital data, image encoding and storage. This process causes different distortions, including geometric and pixelwise distortions. The former introduces perspective transformation and varies the spatial positions of objects within the image, while the latter changes the pixel intensity or color values within the image.

Proposed Method

The framework of the proposed method is shown in Fig. 4, which contains three components: a template-based watermark embedding module, a ceramic pattern-tailored distortion layer, and a learning-based watermark extraction module. Next, we detail each component below. For the convenience of description, we refer to 'pattern' as 'image' in this section.

Template-based Watermark Embedding

The module consists of watermark coding and template embedding, where the former maps the watermark to an invisible template, while the latter adaptively adds the template to the cover image to create the watermarked image.

Watermark coding. Given a watermark message $M = \{0,1\}$ with length $L = 28$ and the cover image $I_o \in \mathbb{R}^{C\times H\times W}$, we aim to generate a watermarked image that contains watermark information while maintaining the visual quality of it as much as possible. Specifically, we first

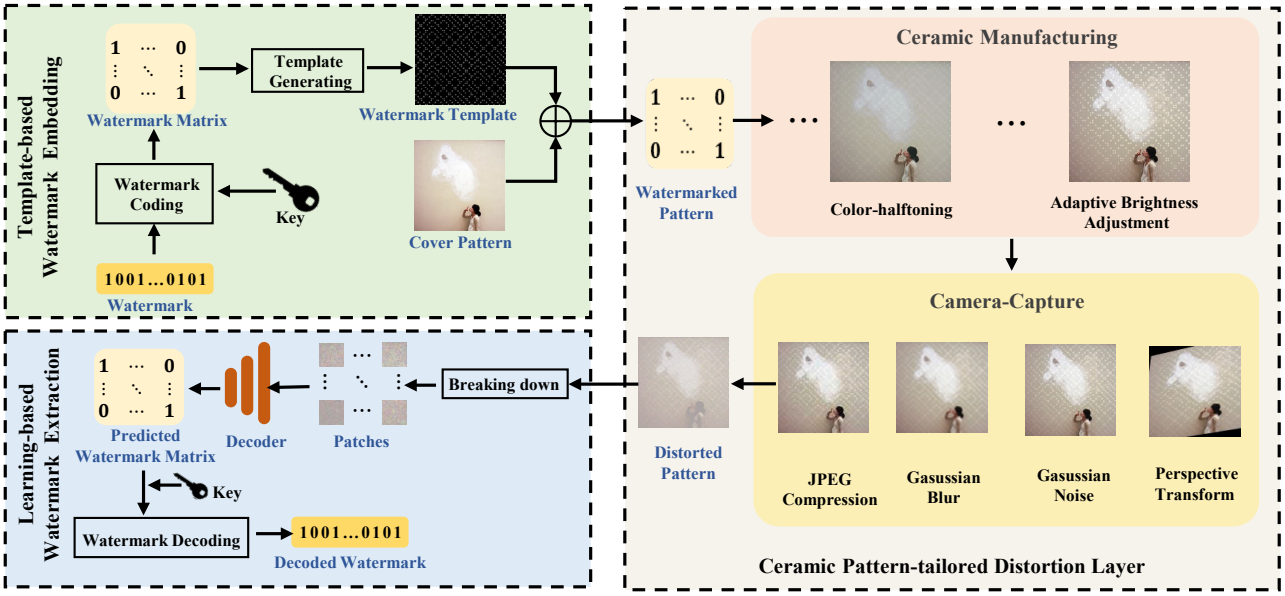

Figure 4: Framework of the proposed method, which consists of three parts: a template-based watermark embedding module to embed watermark into cover images to generate watermarked images, a ceramic pattern-tailored distortion layer that approximates the space resulting from the distortion in ceramic pattern decal-shooting scenario, and a learning-based watermark extraction module to train a deep decoder to extract the hidden data from the distorted watermarked images.

split the watermark message into groups of seven and encode it by Bose Chaudhuri Hocquenghem (BCH) (Bose and Ray-Chaudhuri 1960) code to generate a watermark bit vector B with the error correction and check capability.

$$B = \overbrace{B_0^1 B_0^2 \cdots B_0^{15}}^{\text{BCH groups}} \cdots B_n^1 B_n^2 \cdots B_n^{15} \tag{1}$$

where the value of n represents the number of groups, and 15 indicates the message length of groups after encoding by BCH code.

In addition, extracted watermark information may be affected by bubble distortions during the firing ceramic process. To alleviate this effect, we embedded the message *4* times to ensure extraction accuracy, generating a robust message bit vector $B^{'}$. Then, the message vector $B^{'}$ is reshaped to the matrix $B^{''}$ sizes of 15×16. Next, we prepare two shuffled identity matrices P sizes of 15×15 and Q sizes of 16×16 determined by the key $K = \{k_0, k_1, \cdots, k_{15}\}$, where P and Q are linearly independent. Finally, a template watermark message matrix is generated by

$$M^{'} = concatenate(K, PB^{''}Q) \tag{2}$$

Template embedding. Considering the issue of bubble noise, we chose a Gaussian circle as the bit template to ensure visual quality and generate a special watermark template $T \in \mathbb{R}^{C \times H \times W}$ based on the bit $'0/1'$ to embed the message matrix into the cover image, where (H, W) is the resolution of the cover image, C is the number of channels. The illustration is shown in Fig. 5. We also set the hyper-parameters α and β to make the trade-off between visual

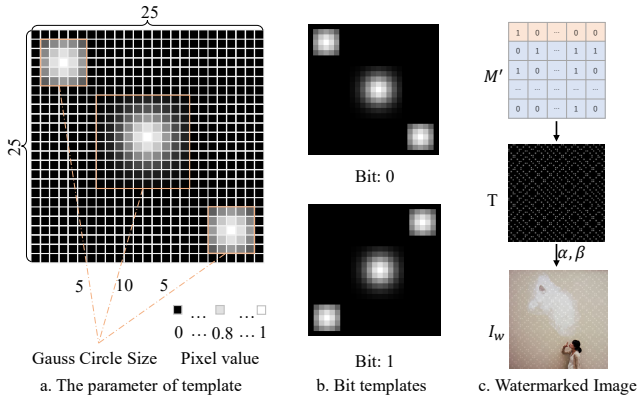

Figure 5: Illustration of the template embedding.

quality and extraction accuracy. The watermarked image can be obtained by

$$I_w = T * \beta * \alpha + I_o * (1 - \alpha) \tag{3}$$

where α and β control the strength of the template.

Ceramic Pattern-tailored Distortion Layer

The distortion layers are crucial for robust watermark extraction. To enhance the robustness of decoder against distortions encountered during the ceramic manufacturing process, we design a ceramic pattern-tailored distortion layer, including the ceramic manufacturing and camera-capture distortion.

Ceramic manufacturing distortion. The main distortion in the ceramic manufacturing process is printing and kiln

firing. The printing process introduces a series of complex physical transformations and noise, where halftone variation is an important factor. Halftone technology (Mese and Vaidyanathan 2002) is a method that simulates continuous tone images using dots of varying sizes and densities, which is widely used in the printing process. We introduce a color-halftoning processing method that combines Gray Component Replacement (GCR) and Bayer matrix ordered dithering to generate the corresponding halftone images. The purpose of Gray Component Replacement is to transfer the gray component from the CMY channels to the K channel in an image, thereby simulating reducing the use of CMY glaze color in the printing process. For each CMYK channel, we first rotate the channel by θ_k degrees. After rotating the image, we apply Bayer matrix dithering to each CMYK channel to avoid moiré patterns. Then, rotate the dithered image back and crop to obtain the final processed image. Finally, merge the processed CMYK channels into a single image. To generate the corresponding halftone image, we reduce it to the size of the corresponding input image, as the image size becomes four times larger during the halftone transformation.

Additionally, we observe that during the firing process, low brightness areas become darker, while high brightness areas may lose color as the glaze can thin out and flake off. Therefore, we incorporate a method into the color adjustment distortion layer that adaptively adjusts brightness differently for various regions based on a brightness threshold to simulate this distortion. First, add a random color offset to each RGB channel sampled uniformly from $[-0.1, 0.1]$. Then use the linear transformation $mx + b$, where $b \sim U[-0.3, 0.3]$ and $m \sim U[0.5, 1.5]$ to adjust contrast and brightness. After this, perform a random proportion linear blend between the RGB image and its grayscale equivalent to adjust saturation. Create a mask to identify pixels with low brightness:

$$M(x,y) = \begin{cases} 1 & \text{if } V(x,y) < t \\ 0 & \text{if } V(x,y) \geq t \end{cases} \tag{4}$$

where t is threshold, we set it as 0.55. Adjust the brightness V component as follows:

$$\begin{aligned} V'(x,y) &= V(x,y) - V(x,y)M(x,y)b \\ &+ V(x,y)(1 - M(x,y))w \end{aligned} \tag{5}$$

where $b \sim [0.2, 0.3]$ and $w \sim [0.1, 0.15]$ are the adjustment ratios for the dark and bright regions, respectively.

Camera-capture distortion. Camera-capture distortion mainly includes spatial and pixelwise, where spatial distortion is caused by the angle of the capture and pixel-wise distortion derived from both camera motion and inaccurate autofocus. For spatial distortion, we perturb the four corner locations of the marker uniformly within a fixed range(up to ±20 pixels) and then perspective transform it to generate a perturbed image. For pixelwise distortion, we randomly sample an angle and generate a straight line blur kernel with a width between 3 and 7 pixels to simulate motion blur, use a Gaussian blur kernel with randomly sampled between 1 and 3 pixels to simulate misfocus. Furthermore, we employ a Gaussian noise model (sampling the standard deviation $\sigma \sim U$ [0, 0.2]) and JPEG compression (Shin and Song 2017) with quality sampled uniformly within [50, 100] to account for noise introduced by camera systems imaging.

Learning-based Watermarking Extraction

Decoder training. The watermarked images are fed into ceramic pattern-tailored distortion layer to generate distorted images I_d. After obtaining the distorted images, we break down it into a sequence of flattened patches $I_d^p \in \mathbb{R}^{C \times N \times L_p \times L_p}$, where (L_p, L_p) is the size of bit template, and $N = HW/L_p^2$ is the number of patches. The patches I_d^p are fed through the decoder to obtain the predicted watermark matrix $\hat{M}_d'$. The loss function of the decoder network can be formulated as:

$$L = \sum ||M^{'} - \hat{M}_w'||^2 + \eta \sum ||M^{'} - \hat{M}_d'||^2 \tag{6}$$

where η is set to 2, $\hat{M}_w'$ denotes the predicted watermark matrix from watermarked image I_w, while $\hat{M}_d'$ is the predicted watermark matrix from distorted image I_d.

Watermark extraction. The predicted key and coded watermark can be obtained by partitioning the predicted watermark matrix $\hat{M}_d'$. Specifically, we split the first row as the predicted Key $\hat{K}$ to obtain shuffled identity matrices P and Q and generate the ordered watermark matrix $\hat{B}''$ using the remaining rows.

$$\hat{B}'' = P^{-1}\hat{M}_d'[1:16, 0:16]Q^{-1} \tag{7}$$

Subsequently, we divide the ordered watermark matrix into *4* groups and then perform voting to obtain the predicted coded watermark $\hat{B}$.

$$\hat{B} = vote(\hat{B}''[:,i], \hat{B}''[:,i+n], \hat{B}''[:,i+2n], \hat{B}''[:,i+3n]) \tag{8}$$

where $i = 0, 1, ..., n-1$ and n represents the number of groups as defined in Equation (2). The final decoded watermark sequence $\hat{M}$ can be obtained by performing BCH decoding on $\hat{B}$.

Experiments

Experimental Settings

The decoder for watermarking extraction takes the ResNet18 architecture (He et al. 2016). To train it, we randomly select 40,000 cover images from the COCO dataset (Lin et al. 2014). Before embedding watermark, we resize them to the resolution of 400 × 400. The length of watermark message is set to 28; each bit in the message is independently sampled from a Bernoulli distribution with probability $\frac{1}{2}$. The size of bit template L_p is set to 25 and the hyperparameter η is set to 2. Unless stated otherwise, α and β are set to 0.25 and 0.4, respectively. In decoder training, color-halftoning (Mese and Vaidyanathan 2002) is a time-consuming operation when processing watermarked images. In real implementation, we use 2000 images to train a U-net (Ronneberger, Fischer, and Brox 2015) to simulate and replace this transformation.

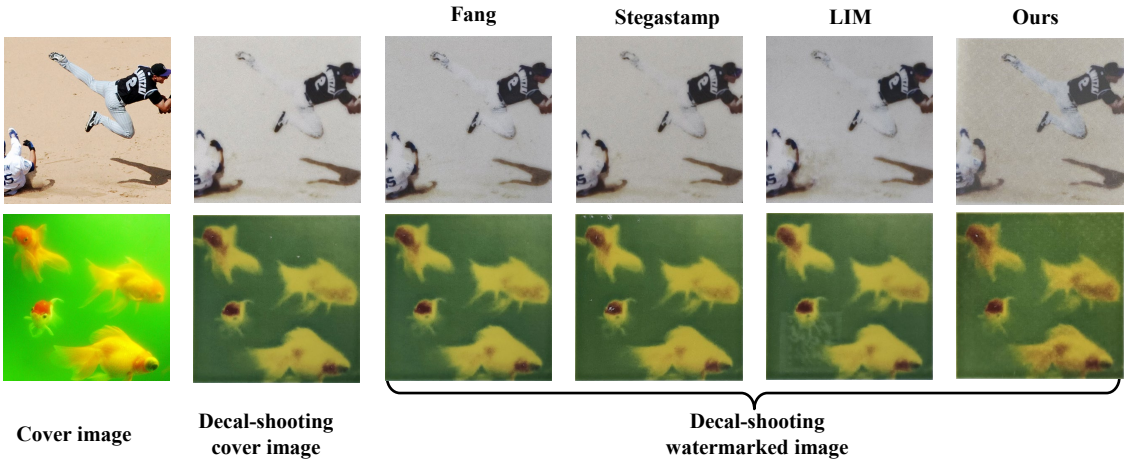

Figure 6: Visual examples of the decal-shooting images by our proposed and comparison methods.

As the first step in exploring the feasibility of embedding watermarks into ceramic artworks' patterns, we evaluate the effectiveness of our method in a simple setting. Specifically, natural images derive from the mirflickr25k (Huiskes and Lew 2008) are adopted as the patterns, and all of them are decaled onto flat tiles. Under this setting, we compare our method with the traditional physical watermarking method (Fang et al. 2018) and the DNN-based ones (Tancik, Mildenhall, and Ng 2020; Jia et al. 2022). Due to the complex and lengthy ceramic manufacturing process, we made every effort to create 10 real porcelain samples for each method. There are two metrics utilized to measure the visual quality of the decal-shooting watermarked image, including Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). The bit accuracy (termed as Bit ACC for short) is adopted to be the indicator to evaluate the robustness of our and the comparison methods. The larger value of PSNR/SSIM and Bit ACC indicate higher visual quality and robustness, respectively.

Metrics	Methods			
	Fang	StegaStamp	LIM	Ours
PSNR (dB)	26.67	22.81	**26.77**	22.79
SSIM	**0.519**	0.460	0.503	0.451
Bit ACC (%)	52.38	80.23	87.50	**97.51**

Table 1: The visual quality of watermarked images and the extraction accuracy of hidden data for different methods, with the best values in bold.

Distance (cm)	Methods			
	Fang	LIM	StegaStamp	Ours
15	52.38	87.50	80.23	**97.51**
20	49.80	84.46	80.39	**94.27**
25	51.06	80.12	78.56	**92.79**
30	47.61	73.69	78.55	**91.65**

Table 2: Bit ACC (%) with different shooting distances.

Comparison against SOTA Methods

To accurately measure the degradation of the decal-shooting image caused by different methods, we calculate the two metrics between the decal-shooting watermarked image and the decal-shooting cover image (i.e., the decaled image on the ceramic artworks without hidden data), and show the results in Table 1. As can be seen, our method achieves favorable visual quality, which is slightly lower than that of the comparison methods (i.e., less than 4 dB lower in PSNR than the best result). In terms of the extraction accuracy of the hidden data, however, the proposed method outperforms all other methods. Specifically, we achieve a 10.01% improvement in Bit ACC over the second-best result. The higher robustness (i.e., 97.51% in Bit ACC) indicates that our method is adequate for use in the majority of real-world ceramic artworks' copyright protection applications. Fig. 6 presents visual examples of the decal-shooting images produced by different methods, with the first and second columns showing the digital cover image and the decal-cover image for clearer comparison. We can observe that there is a significant visual difference between the two different types of cover images, which indicates that the ceramic firing and shooting environment introduces considerable distortions. Compared with other methods, our decal-shooting watermarked images show higher color saturation and reveal almost the same image content and details.

Visual Quality under Different Shooting Settings

In this part, we test the robustness of the proposed method under different camera-capture settings. Specifically, we decal the watermarked images on flat tiles and use mobile phones to capture them in different conditions. For watermark message extraction, we first perspective correct each captured watermarked image and then fit it into the decoder. The variables in the setting include camera shooting distance, shooting angle, and shooting equipment.

For the shooting distances, we slightly change the image

Angle (°)	Fang	LIM	StegaStamp	Ours
Left/Right 15	51.19	79.03	76.60	**86.56**
Left/Right 25	49.31	83.53	77.59	**84.06**
Left/Right 35	49.25	82.42	75.61	**84.37**
Left/Right 45	48.17	77.86	73.80	**82.50**

Table 3: Bit ACC (%) with different shooting devices angles.

Type	Fang	StegaStamp	LIM	Ours
Vivo	52.38	80.23	87.50	**97.51**
iPhone	53.65	80.97	84.42	**95.39**

Table 4: Bit ACC (%) with different shooting devices.

shooting distances of [15,30] cm in the test stage, as shown in Table 2. At all distances, the bit accuracy of our model exceeds 90%, and our model has noticeable advantages over several other models. As the shooting distance increases, the extraction accuracy of the method decreases accordingly.

For the shooting angles, we conducted experiments at different shooting angles to account for potential tilt issues during the shooting process. We first corrected the captured images and then extracted them. The experimental results are presented in Table 3. We found that the tilt angle significantly affects the accuracy of watermark extraction for each method. To further verify the robustness of our method, we conduct experiments with different types of mobile phones and invariable distance fixed at 15 cm. From Table 4, it can be observed that while there are slight variations in bit extraction accuracy using different mobile phones, our message extraction accuracy remains consistent. This indicates that our model has generalization capability across different devices and demonstrates the effectiveness of our strategy.

Ablation Study

The influence of intensity parameters. To choose the most suitable embedding intensity parameters α and β, supplementary experiments were conducted on the same test data, where α represents the embedding intensity of the template, and β indicates the strength of the template, as shown in Table 5. It can be seen that α significantly affects the extraction accuracy and the visual quality of watermarked images. With the increase of α, the PSNR value decreased, but the extraction accuracy increased.

In addition, we visualized the images captured after decaling with different parameters and magnified the bottom right corner of each image to illustrate the differences in the images, as shown in Fig. 7. Based on the results, considering both the visual effect and the accuracy of information extraction, we ultimately fixed parameters α to 0.25 and β to 0.4 in the experiments.

The effectiveness of the proposed ceramic pattern-tailored distortion layer. To further show the advantages of the proposed distortion layer, we train the decoder with-/without the color-halftoning and adaptive brightness adjustment strategy, the results are exhibited in Table 6. As we can see from Table 6, the bit accuracy with color-halftoning

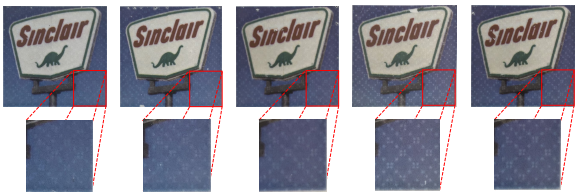

Figure 7: Visual examples of the watermarked images generated by our method with different α and β.

Parameters α	Parameters β	Bit ACC (%)	Message Acc (%)	PSNR
0.15	0.5	86.25	40	23.54
0.15	0.6	91.25	80	22.79
0.25	0.4	97.51	100	22.51
0.25	0.5	96.50	100	21.13
0.25	0.6	95.93	100	19.97

Table 5: Experimental results with different intensity parameters α and β.

Color-halftoning	Brightness adjustment	Bit ACC (%)
×	×	92.65
✓	×	94.06
×	✓	93.90
✓	✓	97.51

Table 6: Effectiveness of color-halftoning and adaptive brightness adjustment strategy.

or adaptive brightness adjustment strategy increases by approximately 1-2%, while with color-halftoning and adaptive brightness adjustment strategy increases by around 5%. This indicates the proposed ceramic pattern-tailored distortion layer effectively simulates decal process, making it more robust for watermark extraction.

Conclusion

In this paper, we explore the feasibility for embedding robust watermarking into patterns to protect the copyright of ceramic artifacts. To achieve this, we develop a novel robust watermarking framework in ceramic pattern decal-shooting scenario, including a handcrafted template-based watermark embedding module, a ceramic pattern-tailored distortion layer and a watermark extraction module. Specifically, we first design a template to embed the copyright information, generating watermarked images. In addition, to guarantee the robustness for decal distortion, we design a ceramic pattern-tailored distortion layer, including color-halftoning, adaptive brightness adjustment and common camera-capture distortion, which can simulate the decal-shooting distortion. Finally, we propose an extraction module to retrieve the embedded message to verify the copyright ownership. Compared to existing methods, our framework exhibits robust performance under real-world decal ceramics scenarios.

Acknowledgements

This work was supported by the National Natural Science Foundation of China under Grants U20B2051, U22B2047, 62450067, 62072114.

References

Akhaee, M. A.; Sahraeian, S. M. E.; and Jin, C. 2011. Blind image watermarking using a sample projection approach. *IEEE Transactions on Information Forensics and Security*, 6(3): 883–893.

Bose, R. C.; and Ray-Chaudhuri, D. K. 1960. On a class of error correcting binary group codes. *Information and control*, 3(1): 68–79.

Fang, H.; Chen, K.; Qiu, Y.; Liu, J.; Xu, K.; Fang, C.; Zhang, W.; and Chang, E.-C. 2023. DeNoL: A Few-Shot-Sample-Based Decoupling Noise Layer for Cross-channel Watermarking Robustness. In *Proceedings of the 31st ACM International Conference on Multimedia*, 7345–7353.

Fang, H.; Jia, Z.; Ma, Z.; Chang, E.-C.; and Zhang, W. 2022. Pimog: An effective screen-shooting noise-layer simulation for deep-learning-based watermarking network. In *Proceedings of the 30th ACM international conference on multimedia*, 2267–2275.

Fang, H.; Zhang, W.; Zhou, H.; Cui, H.; and Yu, N. 2018. Screen-shooting resilient watermarking. *IEEE Transactions on Information Forensics and Security*, 14(6): 1403–1418.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Huiskes, M. J.; and Lew, M. S. 2008. The mir flickr retrieval evaluation. In *Proceedings of the 1st ACM international conference on Multimedia information retrieval*, 39–43.

Jia, J.; Gao, Z.; Zhu, D.; Min, X.; Zhai, G.; and Yang, X. 2022. Learning invisible markers for hidden codes in offline-to-online photography. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2273–2282.

Lin, C.-Y.; Wu, M.; Bloom, J. A.; Cox, I. J.; Miller, M. L.; and Lui, Y. M. 2001. Rotation, scale, and translation resilient watermarking for images. *IEEE Transactions on image processing*, 10(5): 767–782.

Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, 740–755. Springer.

Liu, G.; Si, Y.; Qian, Z.; Zhang, X.; Li, S.; and Peng, W. 2023. WRAP: Watermarking Approach Robust Against Film-coating upon Printed Photographs. In *Proceedings of the 31st ACM International Conference on Multimedia*, 7274–7282.

Mese, M.; and Vaidyanathan, P. 2002. Recent advances in digital halftoning and inverse halftoning methods. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 49(6): 790–805.

Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.

Shin, R.; and Song, D. 2017. Jpeg-resistant adversarial images. In *NIPS 2017 workshop on machine learning and computer security*, volume 1, 8.

Tancik, M.; Mildenhall, B.; and Ng, R. 2020. Stegastamp: Invisible hyperlinks in physical photographs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2117–2126.

Wang, G.; Ma, Z.; Liu, C.; Yang, X.; Fang, H.; Zhang, W.; and Yu, N. 2024. MuST: Robust Image Watermarking for Multi-Source Tracing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 5364–5371.

Wengrowski, E.; and Dana, K. 2019. Light field messaging with deep photographic steganography. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1515–1524.

Zareian, M.; and Tohidypour, H. R. 2014. A novel gain invariant quantization-based watermarking approach. *IEEE Transactions on Information Forensics and Security*, 9(11): 1804–1813.

Zheng, D.; Zhao, J.; and El Saddik, A. 2003. RST-invariant digital image watermarking based on log-polar mapping and phase correlation. *IEEE transactions on circuits and systems for video technology*, 13(8): 753–765.

Zhu, J.; Kaplan, R.; Johnson, J.; and Fei-Fei, L. 2018. Hidden: Hiding data with deep networks. In *Proceedings of the European conference on computer vision (ECCV)*, 657–672.