EG UK Computer Graphics & Visual Computing (2018)
G. Tam and F. Vidal (Editors)

Single Image Watermark Retrieval from 3D Printed Surfaces via Convolutional Neural Networks

Xin Zhang, Qian Wang and Ioannis Ivrissimtzis[1]

[1]Department of Computer Science, Durham University, UK

Abstract

In this paper we propose and analyse a method for watermarking 3D printed objects, concentrating on the watermark retrieval problem. The method embeds the watermark in a planar region of the 3D printed object in the form of small semi-spherical or cubic bumps arranged at the nodes of a regular grid. The watermark is extracted from a single image of the watermarked planar region through a Convolutional Neural Network. Experiments with 3D printed objects, produced by filaments of various colours, show that in most cases the retrieval method has a high accuracy rate.

CCS Concepts

•***Computing methodologies*** → ***Computer vision;*** *Image manipulation;*

1. Introduction

Additive manufacturing, most commonly known under the name of *3D printing*, is a manufacturing method that creates objects by adding one layer of material on top of the other. One of the advantages of additive manufacturing over more traditional methods is that it can handle geometries which either were impossible to manufacture, or the cost of their manufacturing was prohibitive [Add17].

Watermarking is the embedding of information on a physical object or a digital file. Watermarks can be visible or invisible, robust or fragile and can serve various purposes, ranging from object authentication, copy control and protection against unauthorised alteration, to the dissemination of machine readable information. The latter is an increasingly popular application, especially in the form of QR codes and is the target application for the proposed method for watermarking 3D printed objects.

In this work-in-progress paper, aiming at utilising the capability of 3D printing to create objects of more complex geometry at no additional cost, we propose a 3D surface watermark similar in design to QR codes. The information is embedded on the nodes of a regular grid, one bit per node, using semi-spherical or cubic bumps. The existence of a bump corresponds to value of 1 at that location of the grid and its absence to a value of 0. Here, we restrict ourselves to planar surfaces and leave the generalisation of the technique to other surfaces as future work. See Figure 1 for an example.

The proposed watermarking of 3D printed surfaces has some obvious similarities with digital 3D watermarking. Indeed, the construction of a 3D printed object starts with the creation of a digital file, which in the case of watermarked 3D printed objects is a watermarked 3D file, and which up to 3D printer limitations and 3D printing uncertainties it specifies the printed object completely. However, the two processes could diverge significantly during the retrieval process. In the case of digital 3D watermarking, the input of the retrieval algorithm is a digital file and the watermark is retrieved by an inversion of the embedding process. The challenge in this case is to make the embedding-retrieval cycle robust against a variety of malicious and unintentional attacks which could even include 3D printing and rescanning [MAM15, HKCL15]. On the other hand, in our approach of 3D printed watermark retrieval, there is no need for reconstructing the 3D file. Instead, we treat watermark retrieval as a computer vision problem, aiming at retrieving the watermark from a single photo of the printed object, with no need for using any specialised equipment such as laser scanners.

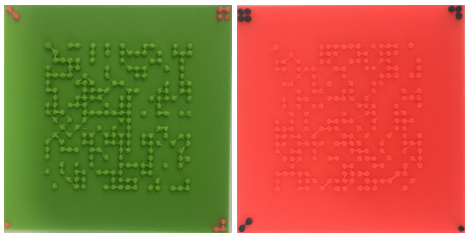

Figure 1: *Examples of 3D printed surface watermarks.*

In this paper we investigate the extent to which deep neural networks can overcome the multitude of challenges that a realistic scenario would pose into the retrieval of such watermarks. In particular, we assume:

DOI: 10.2312/cgvc.20181215

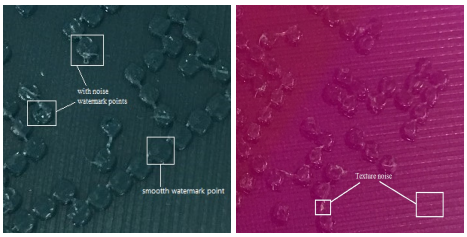

Figure 2: *Left: poor printed watermarks bumps. Right: background artifacts.*

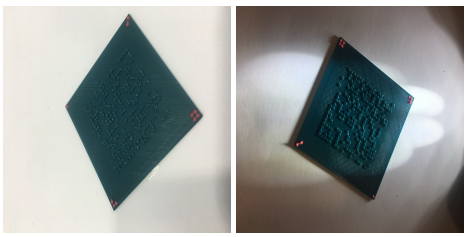

Figure 3: *Watermark retrieval challenges due to illumination. Left: natural light. Right: uneven artificial light.*

1. Use of a single material for the main watermark, meaning that the carrier surface and the bumps on it are of the same colour.
2. Watermark retrieval from a single image from a hand-held phone camera.
3. Use of a low end printer with plastic filaments, meaning that the carrier surface may contain significant noise as well as various infill patterns created during slicing. For that reason, the sizes and shapes of the bumps may vary considerably, even if they are identical in the digital file. See Figure 2.
4. Arbitrary viewpoint and variability of lighting conditions, from natural light to uneven artificial illumination. See Figure 3.

The main contribution of the paper is the development of a convolutional neural network we named CNN-3DW. It takes an RGB image as input and outputs a density map representing the probability of a pixel to be on a watermark bump. The next steps are more straightforward, that is, the registration of the image of the density map, followed by rather simple image processing and analysis operations. In a limitation of our approach, aiming at decoupling the main machine learning component from the subsequent image processing and analysis steps, we opted for a convenience solution into what is a classic image registration problem and we used a second material of distinct colour to mark the four corners of a square region of interest. As a robust convenience solution to the watermark orientation problem, we used four distinct corner symbols consisting of one, two, three and four small dots, respectively.

2. Related Work

Comprehensive reviews of additive manufacturing with academic or industrial focus can be found in [WH12, Ter12]. Currently, there is a multitude of competing additive manufacturing technologies, from laminated object manufacturing to photosensitive polymer curing molding to laser powder sintering molding. As each technique has it sown special characteristics and limitations the use of only one 3D printer to create all our test objects is another limitation of our approach.

2.1. Digital 3D Watermarking

Digital 3D watermarking is an active, well-developed field [OMT02, Bor06, LB11, YPRI17], mostly focusing on boundary surface representations and triangle meshes in particular.

The survivability of the surface watermarks under 3D printing and retrieval by laser scanning has been tested in [MAM15], and surface watermarking algorithms resilient to the 3D printing attack have been developed in [HKCL15]. Notice that the aim of the watermarking algorithm in [HKCL15] is the protection of the digital 3D file and thus, the laser scanning of the 3D printed object followed by surface reconstruction is a required step of the retrieval process. In contrast, as our focus is the embedding of machine readable information on the 3D printed object, we see watermark retrieval as a classic computer vision problem.

2.2. Feature extraction

Convolutional neural network have demonstarted impressive performance on a variety of computer vision tasks. They are considered particularly well suited for feature extraction and classification tasks, AlexNet [KSH12], VGGNet [SZ14] and GoolgeNet [Sze15] being some famous examples. Here, following Zhang *et al.* [ZZC*16] we take the simple but effective approach of training a single column CNN to extract a feature probability density map.

Regarding more traditional image processing the analysis techniques, in an earlier approach to the problem of 3D printed watermark extraction we developed a technique based on local binary patterns (LBP) [OPM00] to recognize the watermark bumps. However, as it is often the case with hand-crafted feature extraction methods, the results did not generalise very well under adverse conditions, in particular, when the background patterns were too prominent, or under extreme uneven illumination, or under unfavourable camera viewpoints.

3. Method

In all of our experiments the watermark was a 20×20 binary matrix. Using filaments of 9 different colours we 3D printed 16 objects in total, each one carrying a different, randomly generated watermark.

3.1. Dataset

From each printed object we captured 15 images of size 3024×4032 under different, arbitrary perspectives, 10 of them under natural light and 5 under extreme artificial illumination. In total we captured 240 images, 80% of which were used for training and 20% for validation. As we did not have a sufficient number of images to train properly the proposed CNN we used data augmentation, which is considered the method of choice in dealing with the problem of small training set. On each training image we applied

three random rotations followed by a random change of its average brightness level, and then randomly sampled 10 patches from each rotated image, creating a training set of 5760 images of size 512×512.

The original 192 training images of size 3024×4032 were annotated by hand, creating binary images which had value 1 at the centres of the watermark bumps and value 0 elsewhere. These binary images were downsampled to size 756×1008 to fit the desired size of the CNN-3DW output. The downsampled binary images were converted to smoother density maps by applying on them one pass of a Gaussian filter with $\sigma = 5$. The filtering process also increased the support of the constructed density maps with the regions of non-zero values roughly corresponding to the watermark bumps, see Figure 4.

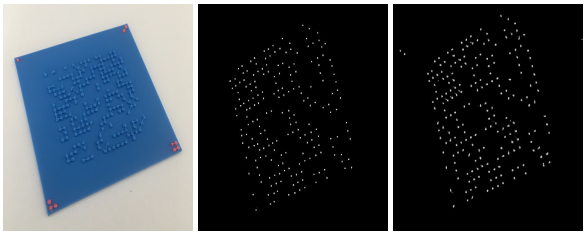

Figure 4: *Left: test image. Middle: ground truth density map obtained by hand annotating the original the image. Right: the density map estimated by CNN-3DW.*

3.2. CNN-3DW architecture

As the various camera views introduce different types of perspective distortion the size of watermark bumps may differ from image to image. However, since the distance between the camera and the printed object is large relative to the size of the object, all bumps in one image are of the same scale size and for that reason we opted for a one column deep neural network to learn the relationship between the original images and the constructed density maps.

The overall structure of the CNN-3DW is showed in Figure 5. There are two layers of max-pooling applied on 2×2 regions, and the activation function is the Rectified Linear unit (ReLu), which generally performs very well in CNNs, see for example [Zei13]. At the top of the CNN-3DW, instead of a fully connected layer we use a convolutional layer of kernel size 1×1. The number of layers, as well as the size and the number of filters in each layer have been chosen through extensive experimentation and cross validation.

We trained the CNN-3DW with Keras built on TensorFlow, using an NVIDIA GeForce GTX TITAN X. We used Adam optimization with a learning rate of $1e-5$ and $1e-6$ after 50 epochs, batch size 2 and momentum 0.9. The average training time was about 7 hours.

3.3. Image registration and matrix retrieval

All models carry four landmarks printed with a different colour and can be easily located using colour segmentation and then used as control points to compute an affine transformation for the image registration. The output density map of the CNN-3DW is thresholded; regularised by applying the same affine transformation on it, and Matlab's *regionprops* function is called to detect its connected regions and obtain estimates of their centroids and their two semi-axes. If the the sum of the two semi-axes is above a threshold we classify that region as a watermark point located at the centroid of that region. See Figure 6.

Finally, to transform the potentially noisy pixel coordinates of the centroids into row and column indices of the watermark matrix we apply k-mean clustering on the x and y coordinates separately, and obtain row and column indices as the indices of the two clusters a centroid belongs, respectively.

	Colour	TPR	SPC	PPV	NPV
1	Green	0.965	0.966	0.967	0.965
2	Dark Brown	0.798	0.863	0.854	0.812
3	Blue	0.964	0.998	0.998	0.966
4	Dark Green	0.662	0.698	0.699	0.661
5	Wooden	0.830	0.840	0.851	0.818
6	Bronze	0.889	0.901	0.900	0.890
7	Luminous Red	0.971	0.981	0.982	0.970
8	Skin	1	1	1	1
9	Green	0.994	1	1	0.993
10	Dark Brown	0.905	0.940	0.937	0.912
11	Blue	1	0.998	0.998	1
12	Dark Green	0.882	0.942	0.939	0.887
13	Wooden	0.862	0.881	0.871	0.872
14	Trans Purple	0.542	0.862	0.756	0.685
15	Luminous Red	1	1	1	1
16	Skin	0.882	0.921	0.0.914	0.0.893
		0.895	**0.897**	**0.895**	**0.885**

Table 1: *Evaluation of the retrieval method.*

4. Experimental results

The test set consisted of three randomly selected images of each printed object, two of them under natural light and one under extreme artificial illumination, for a total of $3 \times 16 = 48$ images.

While during the development of the CNN-3DW network we used Mean Absolute Error and Mean Square Error to quantify the accuracy of the estimated density maps, here we focus on the evaluation of the overall watermark retrieval method and report sensitivity (TPR), specificity (SPC), precision (PPV) and negative predictive value (NPV). The results are summarised in Table 1.

From Table 1 we notice that, generally, the method works well and in conjunction with the use of error correction codes seems to be a viable solution for embedding machine readable information in a physical object. Regarding material suitability, we notice that the results are slightly worse when dark colours are used, such as Dark Brown or Dark Green. We also notice that the algorithm performs worst with the Transparent Purple material, something that could be attributed either to the transparency of the material, or to the use of one only 3D printed object from that material, while most of the other materials were used to print two objects.

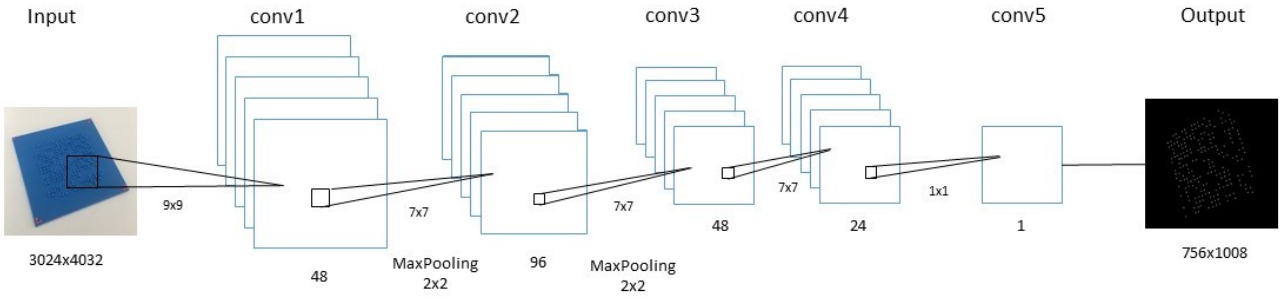

Figure 5: *The architecture of the proposed CNN-3DW.*

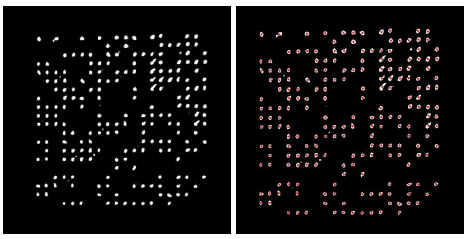

Figure 6: *Left: the regularised CNN-3DW output. Right: the detected centroids of watermark bump regions.*

For the challenges shown in Figure 2 and 3 in particular, our method performs relatively well. For the objects in Figure 2(left) image, the TPR arrives 0.961, while for Figure 2(right), the TPR is 0.860. Under the natural light illumination shown in Figure 3(left), its retrieval accuracy is 0.980, while under the artificial uneven lighting in Figure 3(right), the accuracy drop to 0.812.

5. Conclusion

We presented a method for watermarking 3D printed objects, focusing on watermark retrieval which, in contrast to digital 3D watermarking case, here is a very challenging computer vision problem. Our results show that CNN based computer vision techniques are capable of tackling effectively challenges related to the variability of parameters such as camera view, illumination conditions and printing material.

In future, we will try to extent our methods to 3D printed watermarks lying non-developable surfaces, addressing this way a limitation of a common practice for embedding machine readable information on an object, that is, printing on paper a QR code and sticking it on the object's surface.

References

[Add17] ADDITIVE MANUFACTURING UK: National strategy 2018 - 25, 2017. 1

[Bor06] BORS A. G.: Watermarking mesh-based representations of 3-d objects using local moments. *IEEE Transactions on Image processing 15*, 3 (2006), 687–701. 2

[HKCL15] HOU J.-U., KIM D.-G., CHOI S., LEE H.-K.: 3D Print-Scan Resilient Watermarking Using a Histogram-Based Circular Shift Coding Structure. In *Proc. of the ACM Workshop on Information Hiding and Multimedia Security* (2015), pp. 115–121. 1, 2

[KSH12] KRIZHEVSKY A., SUTSKEVER I., HINTON G. E.: Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (2012), pp. 1097–1105. 2

[LB11] LUO M., BORS A. G.: Surface-preserving robust watermarking of 3D shapes. *IEEE Trans. on Image Processing 20*, 10 (2011), 2813–2826. 2

[MAM15] MACQ B., ALFACE P. R., MONTANOLA M.: Applicability of watermarking for intellectual property rights protection in a 3D printing scenario. In *Proc. of the International Conference on 3D Web Technology* (2015), ACM, pp. 89–95. 1, 2

[OMT02] OHBUCHI R., MUKAIYAMA A., TAKAHASHI S.: A frequency-domain approach to watermarking 3D shapes. *Computer Graphics Forum 21*, 3 (2002), 373–382. 2

[OPM00] OJALA T., PIETIKÄINEN M., MÄENPÄÄ T.: Gray scale and rotation invariant texture classification with local binary patterns. In *ECCV* (2000), Springer, pp. 404–420. 2

[SZ14] SIMONYAN K., ZISSERMAN A.: Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556* (2014). 2

[Sze15] SZEGEDY, C. ET AL.: Going deeper with convolutions. In *IEEE CVPR* (2015). 2

[Ter12] TERRY W.: Additive manufacturing and 3d printing state of the industry. *Annual Worldwide Progress Report, Wohlers Associations* (2012). 2

[WH12] WONG K. V., HERNANDEZ A.: A review of additive manufacturing. *ISRN Mechanical Engineering 2012* (2012). 2

[YPRI17] YANG Y., PINTUS R., RUSHMEIER H., IVRISSIMTZIS I.: A 3d steganalytic algorithm and steganalysis-resistant watermarking. *IEEE transactions on visualization and computer graphics 23*, 2 (2017), 1002–1013. 2

[Zei13] ZEILER M. ET AL.: On rectified linear units for speech processing. In *Acoustics, Speech and Signal Processing (ICASSP)* (2013), IEEE, pp. 3517–3521. 3

[ZZC*16] ZHANG Y., ZHOU D., CHEN S., GAO S., MA Y.: Single-image crowd counting via multi-column convolutional neural network. In *IEEE CVPR* (2016), pp. 589–597. 2