https://doi.org/10.2352/ISSN.2169-4451.2018.34.196

Authentication of 3D Printed Parts using 3D Physical Signatures

Stephen Pollard, Guy Adams, Faisal Azhar and Fraser Dickin; HP Labs; Bristol, United Kingdom

Abstract

We present a flexible workflow for the authentication of 3D printed parts and a series of experiments to show that 3D physical signatures extracted from the surfaces of 3D printed parts are able to robustly and uniquely identify and differentiate otherwise identical printed examples. This forms a useful role within the contexts of track-and-trace, product authentication and anti-counterfeiting. It does not require the product itself to be marked in a specific fashion, thus it does not affect the aesthetics or structural integrity of the printed product.

Introduction

In the world of 3D printed manufacturing it will be important to prove the authenticity of 3D printed parts in order to maintain trust. This will be of particular importance for high value parts being used in critical applications where an approved print process must be maintained. Equally, it will be important to maintain copyright and prevent counterfeit printed parts entering the ecosystem as well as tracing failed parts when problems arise. To this end we are interested in using 3D physical signatures to authenticate 3D printed parts (similar approaches for 2D documents are presented [1]). These are, individual signatures for each printed part based on the random physical structure of the part at the micro scale. This approach has 2 aspects, first it is necessary to identify the approximate (within about 1mm) location and orientation on the printed part from which the signature is to be extracted, and second we must provide a robust and reliable way to capture and compare the physical signatures.

In order to automate the process we propose to define the location of the signature with respect to the CAD model from which the 3D part was printed and identify the location on each instance of the part using automated 3D part alignment and a calibrated robot arm. This paper outlines such a system, but its main focus is a series of experiments showing the utility of 3D physical signatures for 3D part authentication.

Prior Art

There is a long history of using 1D or 2D physical signatures to authenticate printed documents based on the fundamental unclonable intrinsic properties of paper and/or print. One of the earliest reported examples of using the random structure of paper to provide a forensic signature for a document is the FiberFingerprint developed at Escher Laboratories [2]. It was based on a 300byte 1D signal extracted along a piecewise linear path defined by a fixed set of fiducial authentication marks. In 2005 the company Ingenia Technology (www.ingeniatechnology.com) introduced Laser Surface Authentication (LSA™) which uses a 1D laser speckle scanning device with multiple photodetectors to provide a unique fingerprint of paper like surfaces [3]. As the unit scans, the fluctuations from mean intensity of each detector are digitized to form the multi-channel signal that forms the fingerprint code of the surface. The PaperSpeckle method [4] uses off-the-shelf commodity USB microscopes (e.g. the DinoLiteTM AM2011 and the Digital Blue QX5) to image small regions of paper (0.5 mm field of view) at a resolution of 512x384 pixels, leading to an impressively small pixel size of just less than a micron on the paper surface. To aid alignment in the experiments, single ink dots/stains (from a pen) were used for localization purposes.

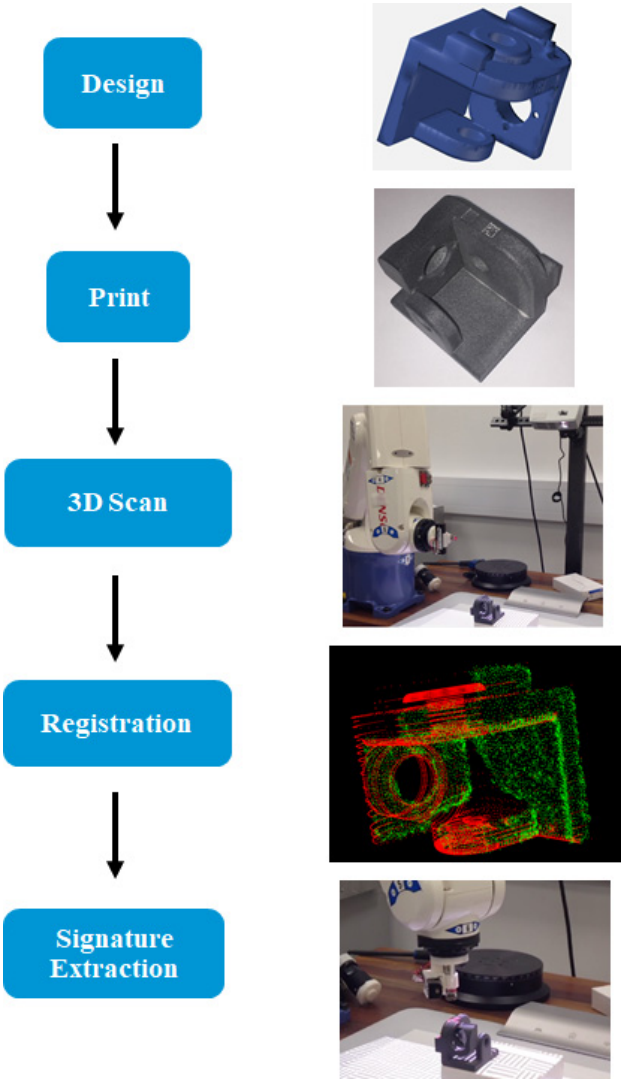

Figure 1. *Shows an outline of the proposed workflow for 3D part physical signature extraction and authentication. See text for details.*

The first example of using printing parasitics as a forensic mark was demonstrated in the Print Signature system of [4]. It used an IntelPlay QX3 low cost digital microscope to capture the forensic mark, termed by the authors a security pattern, which is composed of typically between one and four 1/360 of an inch dots but could in theory be any printed glyph. The

microscope has a resolution around 21000ppi on the paper with the diameter of each dot subtending about 60 pixels. Each glyph was processed to recover N radial components representing the extent of the dot in that direction. A series of papers [5-9] introduce DrCID, a purpose-built document authentication contact imaging device based on a Dyson Relay Lens. Their approach allowed any appropriately modeled glyph, barcode or halftone pattern to act as a forensic mark which was identified in the DrCID image using image registration. Deviation from the model was coded into an authentication signature which provided high levels of statistical robustness.

Similar approaches have been proposed to authenticate 3D objects for example the Fraunhofer Institute in Munich have developed a system for the track and trace of production parts (https://www.ipm.fraunhofer.de/content/dam/ipm/en/PDFs/product-information/PK/IMT/Track-trace-FINGERPRINT-en.pdf) using image based fingerprints recovered from their surface in a controlled production environment.

Authentication Work Flow

Figure 1. shows an outline of a proposed automated 3D part physical signature extraction and authentication workflow. At design time, or shortly afterwards, the location of one or more virtual forensic marks are defined with respect to the CAD model. These are merely locations on the object from which physical signatures are to be extracted and authenticated. There is no need to print any special 2D marker or 3D relief, however, it might sometimes be desirable to provide additional serialization information to assist other aspects of the workflow including primary identification of the printed part. This could be in the form of a standard barcode or RFID tag, either introduced as part of the printing process or after printing is complete. The availability of serialization information can simplify the authentication process as the signature can be associated, in an online database, with the serialization data, making the authentication process a one to one verification rather than a many to one identification process, thus reducing computation effort and improving statistical robustness.

The part is then printed. In our case using an HP Multi Jet Fusion (MJF) 4200 printer [10]. Nylon sintered parts from such a device are mechanically accurate as a result of the fine nylon powder (PA 12 with a 20-60μm particle size) and good quality but only available in a single (greyish black) tone. They can be dyed to produce a more desirable finish.

The remaining 3 stages in Figure 1. illustrate an automated part inspection method of a 3D physical signature extraction process suitable for either the enrollment or authentication stages of the overall system. In order to verify their unique identity signatures must be extracted from 3D parts shortly after printing and stored in a database. This is the recruitment phase of the process. Subsequently, to authenticate the part, the signature must be extracted again and compared to the original to ensure that the physical signatures agree.

In order to automate the signature extraction process we propose the use of a 3D scanner (in this case an HP 3D PRO S3 [11]) and a calibrated robot arm (here we use a Denso VP-6242 industrial robot) incorporating a higher resolution capture device to recover the actual signature. In our experiments we have achieved a closed loop accuracy of about 0.25mm RMS error between the camera system and the robot arm. Printed 3D parts placed in the robot's workspace are scanned to recover triangulated 3D point clouds which can be compared with the original CAD model to register one against the other. The registration phase of Figure 1. shows how the red CAD model is aligned to the green triangulated point cloud. This identifies both the location of the printed part with respect to the robot and more importantly the location upon the part of the virtual forensic mark. This allows the arm to be aligned normal to the region containing the virtual mark for signature extraction as illustrated in the last stage of the process in Figure 1. In the absence of a suitable miniature device with the required 3D capture characteristics for 3D signature extraction, the prototype system shown in Figure 1. merely incorporates a laser device with cross shaped holographic element to indicate alignment.

3D Physical Signature

In this paper we are exploring the utility of a 3D physical signature for part authentication. The signature is based on one proposed in [12] using relocatable features points extracted at a single scale. Here we apply the approach to depth images produced by an Alicona InfiniteFocusSL [13] using a 5X objective lens. This has X and Y spatial resolution of 1.75μm for 2040x2040 samples covering a field of view of 3.6x3.6mm. It has a depth resolution of 0.4μm. While it is infeasible for the Alicona device to form part of our workflow, as it is too bulky to be operated on a robot arm, it is useful for baseline performance experimentation.

Example 3D data is shown in Figure 2. for a region of a Multi Jet Fusion raw part. The surface of this region was intentionally printed to have a random relief with sub 1mm regions printed at different layers less than 0.5mm apart.

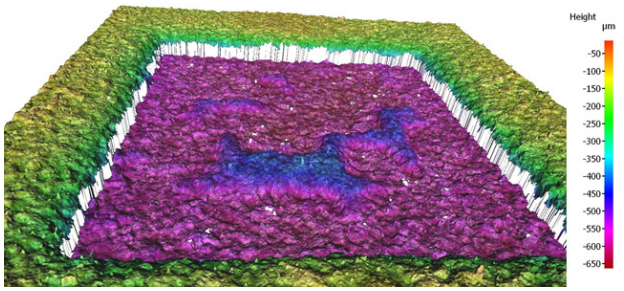

Figure 2. *Shows Alicona 3D data of an MJF part with random relief in the countersunk region to the center.*

Single Scale Features

Single scale features, SSFs, are similar in structure to the scale invariant features of SIFT [14] except in the very important regard that they are extracted at a single fixed scale.

Features are first detected in a difference of Gaussian (DoG) image constructed from the original image or depth map by Gaussian smoothing at a given scale (σ) and at twice that scale (2σ) and then taking the difference. Features are indicated at maxima and minima of this difference image provided their absolute value is above a threshold and they satisfy a Hessian ratio test [15]. The Hessian is approximated on the DoG image at a sampling scale, σ, that is consistent with the scale of the first Gaussian. The Hessian ratio test removes features that are elongated and likely to derive from edge like structures rather than isolated features.

SIFT like descriptors [16] are generated for each detected feature point. These code the relative orientation of the depth image around the detected feature. Orientation and gradients are computed over a 16 x 16 grid centered on the detected feature and scaled by a factor s = max(σ/2, 1.0). The grid itself

can be oriented around the detection point. In a first stage a global orientation histogram is built with contributions weighted according to the gradient of the edge and radial position (according to a Gaussian with standard deviation 4.0). Significant peaks in the histogram (above 80% of the maximum peak) correspond to possible oriented feature descriptions.

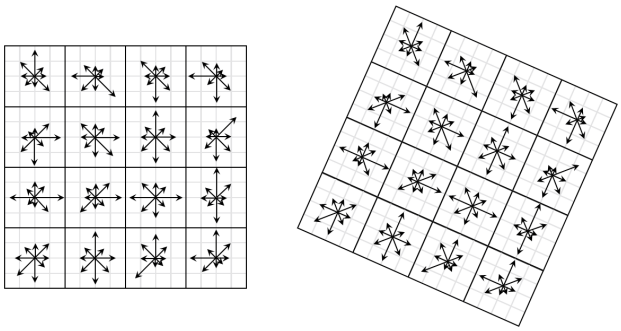

Figure 3. *Illustrates the geometrical form of the 128 elements of a SIFT like feature descriptor. It is shown unrotated on the left and rotated according to the global orientation histogram on the right.*

For each peak angle a second 16 × 16 (s scaled) grid of orientations and gradients is recovered according to the recovered angle (see Figure 3. for an illustration and Figure 4 for an example of grid outlines of detected depth image features). This grid is coded as a 128-element vector comprising 16 histograms each of eight members constructed from contiguous 4 × 4 windows of the grid. With contributions to each histogram again weighted according to the gradient of the edge and radial position with respect to the detected feature location (although this time according to a Gaussian with standard deviation of 8.0). The overall feature vector is normalized to a length of 1.0.

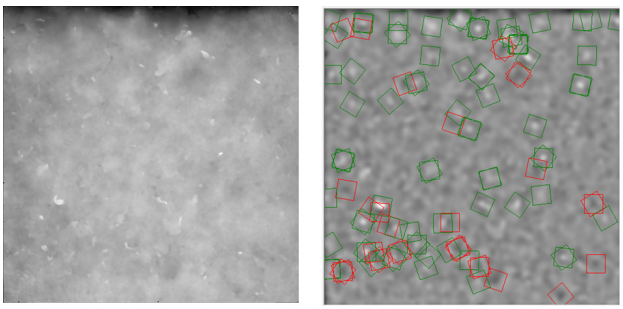

Figure 4. *Shows, on the left, an example 2040 x 2040 depth map extracted from a planar region of an MJF part and on the right, the respective DoG image for σ = 16 pixels along with detected feature points showing the scale and orientation of the (possibly multiple) 16 x 16 grid at each location. The polarity of each feature point is indicated by the color of the square: green for bright features and red for dark ones.*

Feature Signatures

The goal of a feature signature is to be a compact verifiable description that is unique to a specific region of a specific printed part. Rather than use all the features in a given region we identify the top N using a non-maximal suppression scheme with a radius R. That is, features are sorted according to their absolute DoG value and selected in turn eliminating other features within the suppression radius. This process ensures that the signature features are distributed throughout the sampled enrolment region and result in a signature that is robust in terms of both the repeatability of the features and the alignment of the test region during subsequent authentication.

The N features then act as a combined relocatable feature signature. Relocatable in the sense that when presented with the same (or similar) region of the part an overlapping set of new features can be recovered for signature authentication.

Signature Authentication

Given a stored signature the process of 3D part authentication proceeds in a similar manner to the signature extraction. In this case however we extract an increased number (10xN) of comparison features. This allows a significant increase in the robustness of the signature comparison process as it does not require the ordinality of the feature strength to be strongly preserved between the signature enrolment and the subsequent authentication. It also allows for the case where the region used for signature enrolment and authentication to not align perfectly (as will be the case in practice). For example, when the signature and authentication test areas do not overlap perfectly it is possible that one of them could include a number of very strong features that are not present in the other. This would strongly bias the feature selection process if the asymmetric approach of having many more authentication features was not followed.

The process of signature comparison proceeds in a two-stage process of first local and then global feature similarity selection. For each feature in the signature its similarity to features in the authentication test region is computed based on the Euclidian distance between their descriptors. Only the nearest neighbors in terms of descriptor similarity are considered for global feature comparisons. Each signature feature can entertain multiple possible feature matches provided they have a similar (Euclidian) distance metric to the nearest neighbor. That is, all feature matches for a specific feature in the signature must satisfy

$$\alpha F_i < F_{min} \qquad (1)$$

where F_i is the Euclidian distance between feature descriptors, F_{min} is the minimum such distance for the given signature feature and α is a fixed threshold less than 1.0.

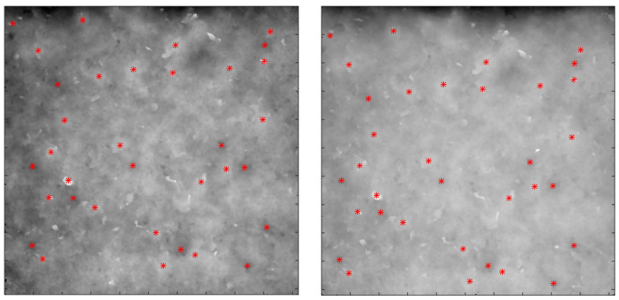

Figure 5. *Shows a signature comprising 50 feature points extracted from the depth image on the left, relocated on a subsequent capture of a similar region on the right. In this case the SSD of the best 10 matching features included in this set was 0.0384 (signature distance 0.196).*

Once the local feature pairings have been recruited the best globally consistent set of feature pairings is constructed. Global consistency is determined by a rough adherence to a 2D Affine transformation model [17]. That is, the largest set of signature feature matches is sought for which a single Affine transformation model will bring them into correspondence with

allowed local feature matches in the authentication set. The latter is defined by a second Euclidian distance constraint but this time in the spatial domain. The distance between Affine transformed signature features and their locally constrained nearest neighbors must be less than T_D.

Global consistency is achieved in a robust and efficient manor using a modified Random Sampling and Consensus (RANSAC) approach [19]. Each putative affine transform is computed from a triple of local feature pairings, where the focus of the triple is considered in turn from the set of locally consistent matches that are unique (no other local feature matches for that signature feature satisfy the constraint in Equation 1.). The other two members of the triple are sought from matches of signature features which are amongst the n nearest spatial neighbors of the of the focus feature and for which the pairwise distances in the signature and authentication spaces satisfy a scale preserving constraint

$$S_{min} < D_s/D_a < S_{max} \quad (2)$$

where D_s is a spatial distance between a pair of features in the signature space and D_a is the corresponding distance between the pair of matching features in the authentication space. S_{min} and S_{max} are a (reciprocal) pair of scale factors close to 1.0.

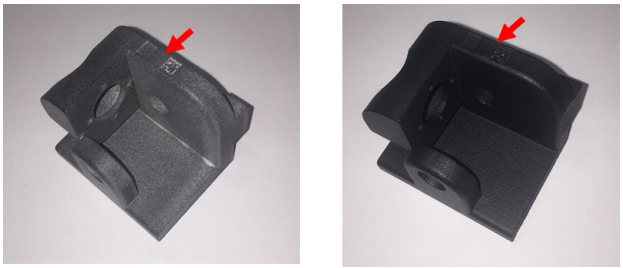

Figure 6. *Shows 2 versions of the same 3D printed part. The one on the left is a raw print while the one on the right has been dyed black to give it an improved finish. The arrow indicates the region used for experimentation.*

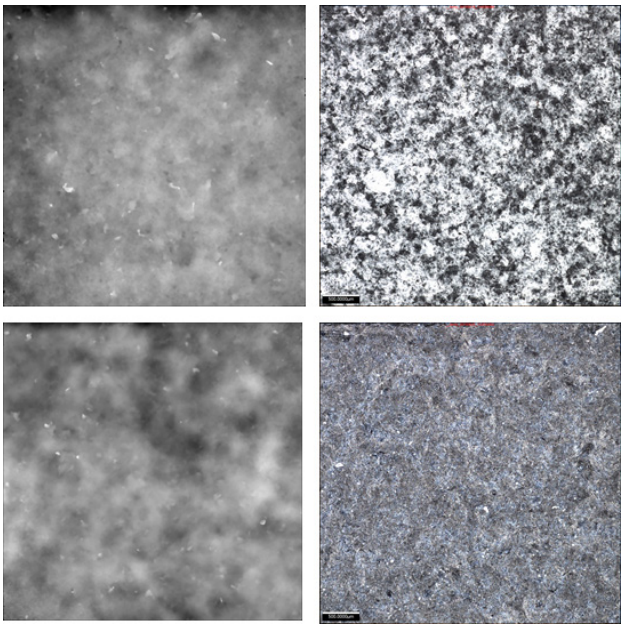

Figure 7. *Shows surface detail from 2 versions of the same 3D printed part. At the top is a raw print while the bottom one has been dyed black. In each case the left-hand side shows a depth map from the Alicona while the right shows the same resolution texture data of the same region.*

The final global authentication distance is calculated based on the sum of the square distances (SSD) of the M smallest feature distances from the set defined by the best Affine consistent global match. The square root of the SSE gives a Euclidian signature distance between the M best matches.

Experiments

We have printed 2 batches of 12 identical 3D parts. One set have a raw finish straight from the printer while the others have been dyed black to give them a better surface visual apperance.

Figure 6. shows examples from each batch while Figure 7. shows depth and texture detail from the Alicona for the regions indicated in Figure 6. Notice that, as expected, the depth data is more consistent than the texture between the raw and dyed versions of the part. We have captured the same/similar regions from all the parts multiple times with and without significant rotation (90 degrees). In a first experiment we consider signature differences for true and false signature comparisons for the raw parts as we alter the σ scale parameter of the DoG filter. All other parameters of the signature capture and comparison remain fixed as shown in Table 1.

Table 1. Experimental Parameter Settings

N	50
M	10
α	0.7
T_D	2.0 (pixels)
S_{min}	0.8
S_{max}	1.2

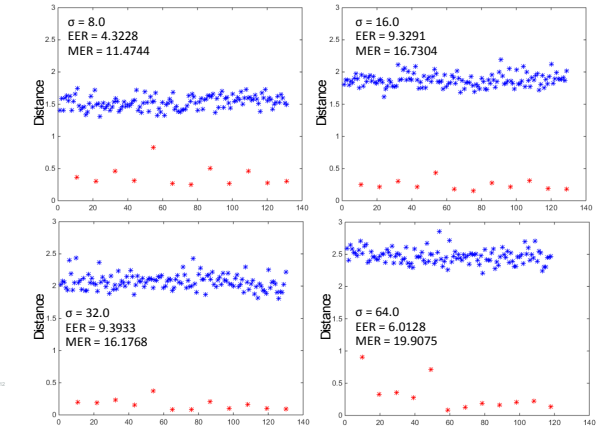

Figure 8. *Signature distance (the square root of the SSD between M=10 best consistent features) plots showing all false comparisons (blue stars) and the few correct comparisons (red stars) as the σ scale parameter of the DoG filter is varied between 8.0 and 64.0. Also shown are approximate error rates between the distributions (see text for details).*

Figure 8. shows signature distance plots for both true and false part comparisons. It can be seen that the distance between the two distributions grows as the scale of the signature features increases. One way to measure the statistical distance between two distributions is in terms of equal error rate (EER). That is the probability at the point where the chances of a false positive and false negative are equal. For Gaussian distributions the EER corresponds to a standard score (or Z-score) of

$$Z_E = \frac{|\mu_2 - \mu_1|}{\sigma_1 + \sigma_2} \quad (3)$$

where μ_1, μ_2 and σ_1, σ_2 are the means and standard deviations of the two Gaussian distributions under consideration. In this case the EER is given by the error function, *erf*, of the normal distribution

$$EER = \frac{1}{2} - \frac{1}{2} erf\left(\frac{Z_E}{\sqrt{2}}\right) \quad (4)$$

For example, the probability of a Z-score of 10 is small indeed at 7.6×10^{-24}. While our distributions are not necessarily Gaussian (especially those for the few positive examples) this provides a useful indication of the statistical separation of the distributions. However, for forensic authentication we are more interested in preventing false positives than allowing the occasional false negative. The former is an indication that a system has been spoofed and a counterfeit has been accepted as a valid example. False negatives on the other hand can be a result of user or equipment error. Thus, the mean error rate (MER) of all positive examples relative to the distribution of false matches can provide a more useful estimate of the system performance. Also, the distribution of false comparison distances follows more closely the Gaussian form, making the estimate more reliable.

While statistical reliability is crucial it is equally important to consider the physical difficulty in spoofing a physical signature which largely comes down to the physical feature size of the elements contributing to it. In Table 2. We show the relationship between the σ scale parameter and both the element size in a feature and the overall extent of the feature itself. For the further experiments we fix σ=32.0 as this provides a good compromise between physical feature size and statistical robustness.

Table 2. Scale of Feature Detector

σ	8.0	16.0	32.0	64.0
Element (μm)	7.0	14.0	28.0	56.0
Feature (μm)	56.0	112.0	224.0	448.0

Figure 9. presents signature distance plots for two more cases. One for the same parts shown in Figure 8. but where the part has been rotated though 90 degrees between signature capture and authentication. And another, for the second batch of parts that have been dyed. Results for the rotated data are very similar to the unrotated data reported in Figure 8. Results for Dyed data on the other hand show a marked improvement due to the reduced difference and variability in the signature distances. This is due to the reduction of the specular reflections of the raw parts that resulted from retained white nylon powder that was not fully removed during cleaning.

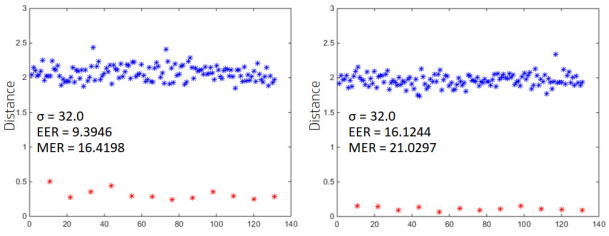

Figure 9. *Signature distance plots showing false comparisons (blue stars) and correct comparisons (red stars) for 90 degree rotated raw parts on the left and dyed parts on the right.*

Finally, in Figure 10. we show comparative data for image/texture based signatures extracted from the same regions as the depth-map based signatures presented previously. The 3 graphs mirror those for σ=32.0 presented in Figures 8. and 9. While qualitatively similar to the results achieved for 3D data we see an improvement in EER for the unrotated case of the raw printed part. MER on the other hand is significantly better for all three cases. This supports the use of high resolution 2D imaging alone as a basis for 3D part authentication.

Conclusions

We have presented a workflow for the authentication of individual 3D printed parts. It has the flexibility to extract physical signatures from arbitrary parts using the CAD model and a calibrated robot arm to align a high-resolution imaging device against a virtual forensic mark defined with respect to the CAD model. In a series of experiments, we have shown that 3D physical signatures based on high-resolution depth information are able to provide very high levels of statistical robustness in their ability to discriminate individual instances of a printed part. However, it is also possible to use 2D physical signatures based on high resolution imaging to achieve similar, if not better, statistical performance. 3D signatures have the advantage that they are coding physical structure rather than just appearance. This has the twin potential benefits of being more difficult to clone and more robust to physical alterations.

However, our experiments were based on an expensive bulky laboratory measurement device and considerable miniaturization and cost reduction will be required to develop a practical 3D physical signature authentication system.

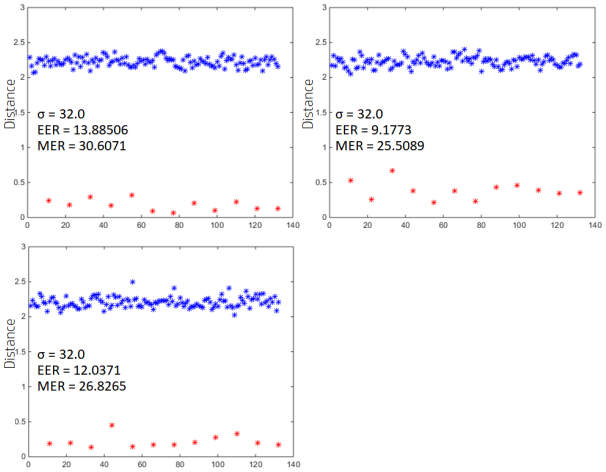

Figure 10. *Signature distance plots showing false comparisons (blue stars) and correct comparisons (red stars) for texture data. Top left is for raw parts, top right is for raw parts rotated 90 degrees and bottom is for dyed parts.*

References

[1] S. Pollard, G. Adams & S. Simske, Forensic Identification of Printed Documents, in Handbook of Digital Forensics of Multimedia Data and Devices, Eds Ho & Li, (Wiley-Blackwell 2015) pg. 442.

[2] E. Metois, P. Yarin, N. Salzman & J.R. Smith, FiberFingerprint identification, Proc. 3rd Workshop on Automatic Identification, pg. 147. (2002).

[3] J.D.R. Buchanan, R.P. Cowburn, A.V. Jausovec, D. Petit, P. Seem, G. Xiong, D. Atkinson, K. Fenton, D.A. Allwood & M.T. Bryan, Forgery: Fingerprinting documents and packaging, Nature, 436 (2005).

[4] B. Zhu, J. Wu, & M.S. Kankanhalli, Print Signatures for Document Authentication, Proc. 10th ACM Conf. on Computer and Communications Security, pg. 145. (2003).
[5] G. Adams, Handheld Dyson Relay Lens for Anti-Counterfeiting, Proc. IEEE Intl. Conf. Imaging Systems and Techniques, pg. 273. (2010).
[6] S. Simske & G. Adams, High-resolution glyph-inspection based security system, Proc. IEEE Intl. Conf. Acoustics Speech and Signal Processing, pg. 1794. (2010).
[7] S. Pollard, G. Adams & S. Simske, Resolving distortion between linear and area sensors for forensic print inspection, Proc. IEEE Intl. Conf. Image Processing, ps. 1001. (2010).
[8] S. Pollard, S. Simske & G. Adams, Model based print signature profile extraction for forensic analysis of individual text glyphs, Proc. IEEE Intl. Workshop on Information Forensics and Security, pg. 1. (2010).
[9] S. Pollard, S. Simske & G. Adams, Print Biometrics: Recovering Forensic Signatures from Halftone Images, Proc. Intl. Conf. on Pattern Recognition, pg. 1651. (2012).
[10] www8.hp.com/us/en/printers/3d-printers.html
[11] https://www8.hp.com/us/en/campaign/3Dscanner/overview.html
[12] S. Pollard, S. Simske & G. Adams, Signature authentications based on features, WO2017176273A1, (2017).
[13] https://www.alicona.com/products/infinitefocussl/
[14] D.G. Lowe, Object recognition from local scale-invariant features, Proc. International Conference on Computer Vision. 2. pg. 1150. (1999).
[15] D.G. Lowe, Distinctive Image Features from Scale-Invariant Keypoints, International Journal of Computer Vision. 60:2. pg. 91. (2004).
[16] http://www.vlfeat.org/api/sift.html
[17] G. Wolberg, Digital Image Warping, (IEEE Computer Society Press Los Alamitos, CA, USA,1994).
[18] M.A. Fischler & R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography, Communications of the ACM,24:6, pg. 381. (1981).

Author Biography

Stephen Pollard is a Distinguished Technologist in the Print Adjacencies and 3D Lab of HP Labs. He has been with the company for over 25 years working on imaging technologies associated with, digital photography, document capture and copying, 3D scanning and document and product authentication. He attained a BSc in Computer Science from the University of Newcastle upon Tyne in 1981 and a PhD in Computer Vision from the AI Vision Research Unit at the University of Sheffield in 1985.

.