Remote Extraction of Latent Fingerprints (RELF)

1st Matthew McGuigan
Computer Science Department
University of Exeter
Exeter, UK
M.McGuigan@exeter.ac.uk

2nd Jacqueline Christmas
Computer Science Department
University of Exeter
Exeter, UK
J.T.Christmas@exeter.ac.uk

***Abstract*—Latent fingerprints are the kind left on objects after direct contact with a person's finger, often unwittingly at crime scenes. Most current techniques for extracting these types of fingerprint are invasive and involve contaminating the fingerprint with chemicals which often renders the fingerprint unusable for further forensic testing. We propose a novel and robust method for extracting latent fingerprints from surfaces without the addition of contaminants or chemicals to the evidence. We show our technique works on notoriously difficult to image surfaces, using off-the-shelf cameras and statistical analysis. In particular, we extract images of latent fingerprints from surfaces which are transparent, curved and specular such as glass lightbulbs and jars, which are challenging due to the curvature of the surface. Our method produces results comparable to more invasive methods and leaves the fingerprint sample unaffected for further forensic analysis. Our technique uses machine learning to identify partial fingerprints between successive images and mosaics them.**

***Index Terms*—Contactless fingerprint extraction, neural network**

I. Introduction

IT has been standard practice to use fingerprints as evidence for decades in criminal convictions and information security. There are three types of fingerprint used in biometrics: latent, patent and plastic [1]. Latent fingerprints are almost invisible, formed by a dielectric residue left behind from the fingerprint ridges containing water with various salts and organic compounds [2]. Extracting latent prints is further complicated since they may be found on complex curved surfaces. Patent fingerprints are easily visible to the naked eye and are formed when the finger is coated in ink or another similar substance then pressed onto a surface. Plastic prints are three dimensional impressions formed when the finger is pressed into a malleable surface such as wax, paint or soap. In this paper we will be dealing with latent prints.

In spite of latent fingerprint extraction being a long established process, invasive techniques are vulnerable to improper collection methods which may cause a loss of information. Latent prints are often enhanced physically for photographs by adding a material which involves 'dusting' the scene in the expectation that the powder will become fixed to the residue left behind, and hence become much more visible in any further imaging. This type of chemical processing may degrade or contaminate the evidence, preventing additional forensic testing [3]. Our method is simple, fast and requires only an off-the-shelf camera and remote flash. Fingerprints are comprised of ridges which may terminate or form bifurcations as well as a variety of other distinctive features formed in the foetus from the fifth month of pregnancy [4]. These features are known as minutiae and usually appear in the fingerprint in unique combinations resulting in one persons fingerprint being clearly discernible from another. Latent fingerprint features are more difficult to match than those of patent fingerprints, and are more susceptible to scrutiny in courtroom arguments [4]. This is due to the non-ideal surfaces (often curved and specular) where latent prints are often found. Specular reflections only reflect along a given direction and a diffuse reflection which will reflect roughly the same intensity in all directions. Specular surfaces complicate imaging forensically because their reflectance is highly direction dependent. For our problem this means the lighting direction has to be just right in order to interact with the specular surface and obtain any meaningful photographic data. Our technique aims to assist in developing a simple and standard non-invasive novel pipeline to extract fingerprints in these circumstances so they may become a less vulnerable form of forensic evidence. Given the scale of fingerprint ridges and the fact that the quality of images impact the effectiveness of fingerprint feature point extraction [5], we work with high resolution input images in order to preserve the finer details.

A. Latent fingerprints and surface curvature

The local shape of a point on a surface is determined by the two principal curvatures (eigenvalues of the shape operator at this point), κ_1 and κ_2. One principal curvature defines the rate of maximum bending and its corresponding tangential direction on the surface, while the other defines the rate and corresponding tangential direction of minimum bending. We will not focus too much on the mathematics of these principal curvatures, but it is important to acknowledge their significance and effects on imaging different surfaces. Two useful quantities we consider are Gaussian and mean curvatures, G and M. As one would expect, the curvature (Gaussian and mean) of a planar surface is zero. However, the mean curvature of a cylindrical surface is greater than zero, while the Gaussian curvature is still zero. This is because the Gaussian curvature of a surface is multiplicative (as shown in (1) and mean curvature is additive (as shown in (2)).

$$G = \kappa_1 \kappa_2 \tag{1}$$

978-1-7281-6926-2/20/$31.00 ©2020 IEEE

$$M = (\kappa_1 + \kappa_2)/2 \tag{2}$$

This means that if any of the principal curvatures, κ_1 and κ_2, are equal to zero then so will the overall Gaussian curvature, G. For this reason the mean curvature is more significant for the purposes of extracting fingerprints. We examine the effects of different principal curvatures and the proposed method's performance, and note how imaging becomes more difficult as mean curvature increases. We extract prints from specular planar surfaces (such as windows) as shown in Fig. 1, with both principal curvatures $\kappa_1 = \kappa_2 = 0$, and hence a mean curvature of zero. Imaging a flat specular surface where both principal curvatures are zero is comparatively easier than when one or both are non-zero (see results in Fig. 7c). We also extract fingerprints from specular cylindrical surfaces (such as glass jars) as shown in Fig. 1 with $\kappa_1 = 0$ and $\kappa_2 > 0$, which results in positive mean curvature. Imaging cylindrical surfaces where one principal curvature is non-zero proves more difficult than a flat surface, but is possible to extract fingerprints using the RELF method as shown in Fig .7f. Our method also proves robust enough to work on specular spherical surfaces (such as spherical bulbs) as shown in Fig. 1 with both principal curvatures κ_1 and κ_2 being positive. Despite the surface being much more difficult to image since both principal curvatures are non-zero, the RELF method is still able to extract fingerprints as shown in Fig 7i and Fig. 8f. The mean curvature of a surface is indicative of the ease at which we may extract fingerprints. The surfaces in these examples are also specular, which poses different issues as outlined in I.

B. Related work

There exist several methods using expensive optical equipment for non-destructive extraction of latent fingerprints from curved smooth surfaces that can yield impressive results. Optical coherence tomography has been used to extract latent fingerprints from complex surfaces [6]. The authors were able to extract fingerprints from even poorly reflecting samples where the latent print was unnoticeable under ordinary viewing conditions, and they achieved this without any physically invasive or chemically enhanced processing.

Another optical method utilises the fact that specularly reflected light from dielectrics is partially polarised at a specific range of observing angles [3]. Despite producing effective results, the techniques in [3], [6] and [7] are intended for extracting fingerprints from flat surfaces. However, another optical method has been developed specifically to image curved surfaces [8]. As the authors state, the ability to obtain a non-destructive reconstruction of a fingerprint (or portion of it) in situ from a cylindrical or curved surface is important for the purposes of identifying a person at a crime scene. This non invasive method uses a diffractive optical element based glossmeter (a device usually used to measure magazine print gloss quality). The method utilises a motor-driven rotary table to rotate the object being imaged. A laser beam is focused onto the position of the latent fingerprint and colour-coded

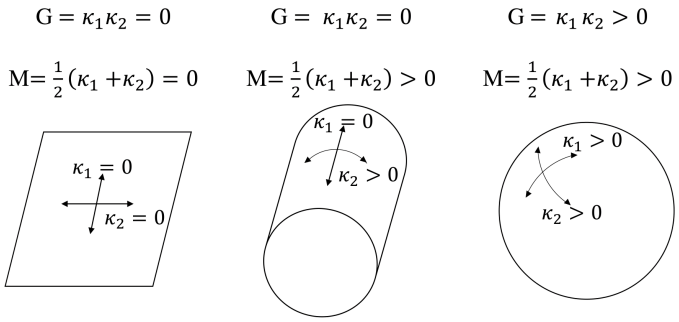

Fig. 1: An outline of the different surface curvatures that the RELF algorithm is equipped to process.

gloss map of the scanned region is obtained, with colour being related to the strength of light reflection. The latent fingerprint is shown in contrast to the background surface due to a difference in reflectivity. The authors state their method may be encounter issues due to colour affecting the image contrast of the fingerprint [8]. It is worth noting that this technique requires lots of equipment and a highly skilled user.

Researchers have used hyperspectral imaging (which constructs a three dimensional data cube consisting of two dimensional images over numerous wavelengths) to obtain fingerprint images. All channels are fused using histogram of oriented gradient information to weigh each of these channels [9]. These non invasive optical methods produce interesting results, but they are often experimental proof-of-concept setups and require a high level of knowledge and skill to operate the equipment. Other invasive and potentially deleterious methods besides dusting include using hardware such as deformable membranes on glass plates and heating glass plates to remove moisture [4]. Clearly, these methods are very invasive which is undesirable in any forensic investigation as they may risk destroying or contaminating the fingerprint sample. Therefore there is clearly a need for a straightforward, non-destructive approach which avoids having to compromise precious forensic evidence. All of these described methods rely on complicated, laboratory based equipment that requires careful calibration. Our proposed method requires only the operation of an off-the-shelf camera and a remote flash.

C. Concepts behind proposed method

It has been noted by forensic investigators for decades that by varying the angle of a torch incident on a surface potentially containing latent fingerprints, it is possible to locate partial or full fingerprints [10]. Our method utilises this basic principle: we illuminate the object suspected to contain a latent fingerprint using multi-light imaging, a technique in which the object is held stationary and a light source illuminates the item from a different direction for each image. Reflectance transformation imaging utilises multi-light imaging (using a lighting dome to avoid the process of manually moving a flash around) to compute surface normals and obtain topographical information about the object being imaged [11]. Reflectance transformation imaging is chiefly used in cultural heritage

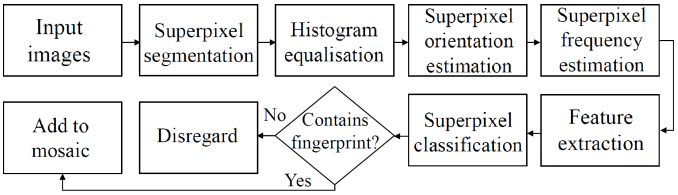

Fig. 2: A flowchart of the RELF algorithm, which extracts features from superpixels for classification then uses these classifications to build a mosaic fingerprint image.

imaging since it is inexpensive and produces high resolution surface topography using off-the-shelf cameras.

We borrow our imaging approach from reflectance transformation imaging, but the similarities end here. We apply machine learning to reveal the hidden latent fingerprints found on objects. These methods have been used before to reveal objects hidden to the human eye, and have been shown to work much better than humans in the dark. Researchers used deep neural networks to identify and improve salient features of dimly lit images [12]. Our technique uses machine learning to identify partial fingerprints between successive images and mosaics them.

II. Methods

The novel pipeline for the proposed RELF method is outlined in Fig. 2. Our approach uses multi-light imaging and an off-the-shelf digital camera to gather images as outlined in II-A. Each of these images is then segmented into superpixels (see Fig. 4) for analysis. We estimate the orientation and frequency of the superpixels in II-C, allowing us to generate a 2D sinusoidal filter from which we can obtain a cross-correlation to see how closely the fingerprint ridges match the peaks and troughs of a sinusoidal function of the same frequency. We extract more features from the superpixel in II-E, reducing the dimensionality of the data thereby allowing us to categorise the superpixel into a binary classification ($1 = fingerprint$ or $0 = nofingerprint$). We will now discuss these steps in more detail.

A. Data acquisition using multi-light imaging

Latent fingerprints viewed under illumination on curved surfaces are only partially visible from any one lighting angle. They also vary in brightness relative to their proximity to the specular reflection on the curved surface. This means we only see small regions of the fingerprint at best in each image. This can be seen in Fig. 3 where the specular reflection is saturated, but the surrounding region contains an eligible portion of the fingerprint. Using the multi-light image collection technique we illuminate the curved glass object to build up, piece by piece, the overall fingerprint image. For each image stack we obtain around 80 images, with the location of the specular reflection being different on the surface because the light changes direction on the lighting dome. At the apex of the multi-light imaging dome an off-the-shelf camera captures an image for each unique lighting direction. Using a lighting

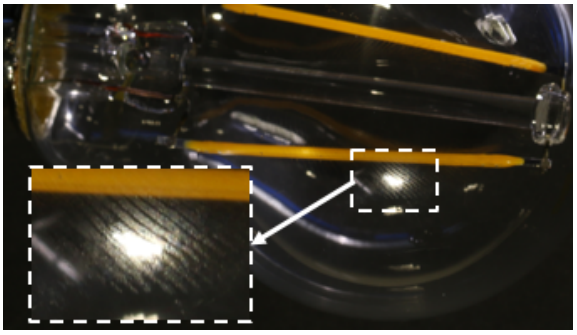

Fig. 3: A light-bulb exhibiting a specular reflection which is saturated and noisy but reveals neighbouring portions of legible fingerprint.

dome is not mandatory but does save time, and we will show RELF performs well on hand-lit image stacks as well as dome-lit image stacks. We will now discuss the proposed algorithm for extracting fingerprints from the image set.

B. Superpixel segmentation

Having obtained the image data, we break down each image into sets of segments called superpixels that collectively cover the entire image. Each individual superpixel in the image contains pixels with similar features such as colour, texture and brightness [13]. Superpixels can be generated by two main categories of algorithms: graph based and gradient ascent [14]. We opt to use a gradient ascent based method known as simple linear iterative clustering (SLIC), which efficiently generates superpixels using k-means clustering [15]. Often specular reflection pixels are grouped together by the SLIC superpixels algorithm due to their similar (saturated and noisy) intensity, we use this to determine whether or not the superpixel requires further processing in the RELF framework. In these instances the entire superpixel is $whited$ out, and yields little to no information. SLIC is straightforward and memory efficient, allows control over the number of superpixels and adheres well to boundaries [14]. We use default parameters for SLIC as follows: compactness is 10, the method is $SLIC0$ as opposed to $SLIC$ so we can adaptively refine compactness after the first iteration and our number of iterations used in the clustering phase of the algorithm is 10.

On a curved and specular surface the intensity of the fingerprint varies drastically with distance to the specular highlight. We use histogram equalisation to even out these disparities. This method is particularly effective for areas with lower local contrast (and further away from the specular reflection) to gain a higher contrast. This increases the clarity of any potential fingerprint ridges and allows us to determine whether any useful portions of the fingerprint are present. Examples of superpixels containing fingerprint data can be seen in Fig. 4. The unprocessed superpixels are shown in (a) - (d) and the corresponding superpixels having undergone histogram equalisation are shown in (e) - (h), where the ridges of the fingerprint portions are much clearer. Now we have increased

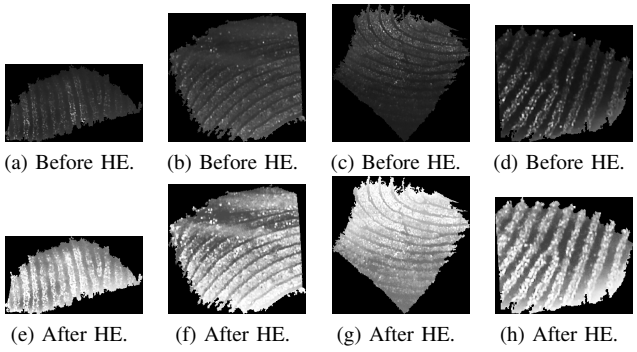

(a) Before HE. (b) Before HE. (c) Before HE. (d) Before HE.

(e) After HE. (f) After HE. (g) After HE. (h) After HE.

Fig. 4: Various input superpixels before and after histogram equalisation (HE) to enhance contrast (a) - (d) the input superpixels containing fingerprint. (e) - (h) the same superpixels after undergoing histogram equalisation.

the clarity of potential fingerprint containing superpixels, we will look at methods for orientation and frequency estimation.

In order to determine whether or not a given superpixel contains a fingerprint portion, we measure several numerical features of the superpixel to build a one dimensional feature vector. We measure fourteen distinct features of the superpixel, and then train a neural network to learn which combination of these features represents a fingerprint. Two of these numerical features, the cross correlation with a 2D sinusoidal filter and the number of fingerprint ridges present in a superpixel, can only be extracted once the Fourier transform is computed. For this reason we will now outline the 2D Fourier analysis performed on each superpixel.

C. Orientation and frequency estimation

We obtain our first two numerical features, the cross correlation and the number of fingerprint ridges present by performing a two-dimensional fast Fourier transform, enabling us to estimate the dominant spatial frequency in the superpixel. We perform Fourier analysis separately on each superpixel, extracting these two features independently for each one (fingerprint ridge frequency is is not assumed as constant due to curved surface). We are then able to measure the orientation of this dominant frequency relative to the horizontal, and calculate the cross-correlation of the superpixel with a filter generated using the underlying predominant spatial frequency of the superpixel. This cross-correlation is our first numerical $feature$ to be input into our machine learning algorithm as outlined in II-E. If a given superpixel indeed contains fingerprint ridges, it will exhibit a high cross correlation with a sinusoid filter of the same frequency as shown in Fig. 5.

We take into consideration that the cross-correlation of a sinusoidal filter with a fingerprint portion may encounter issues due to the fact that the curvature of fingerprint ridges increase towards the centre of the fingerprint [16]. In these central fingerprint sections with high ridge curvature, the assumption of a dominant ridge direction and parallel ridges is not valid since the curvature is too great. This means, in

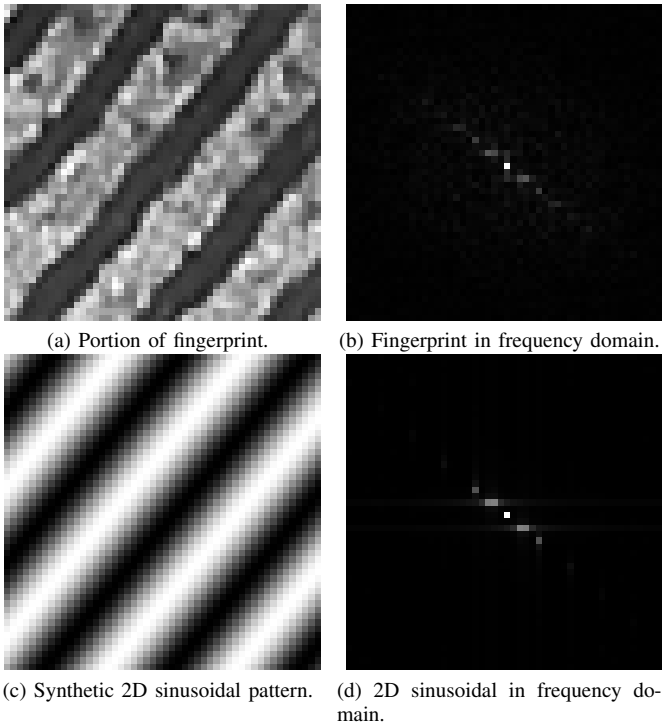

(a) Portion of fingerprint. (b) Fingerprint in frequency domain.

(c) Synthetic 2D sinusoidal pattern. (d) 2D sinusoidal in frequency domain.

Fig. 5: (a) Sample of exemplar fingerprint portion (b) The resulting frequency domain image computed from the Fourier transform of the fingerprint portion. (c) A 2D sinusoidal filter. (d) The resulting frequency domain image after a Fourier transform on the sinusoidal filter.

central fingerprint regions, the correlation could indicate a low similarity with the sinusoidal filter. However, the effects of this issue were found to be minimal if the number superpixels used is sufficiently high. This is because if we increase the number of superpixels (and hence decrease their size) the central regions appear to be approximately less curved.

Indeed, most of the superpixels containing fingerprint ridges (see Fig. 5a) were found to be sufficiently parallel that they closely match the sinusoidal filter (see Fig. 5c). These similarities are also visible in the corresponding frequency spectra of the fingerprint (see Fig. 5b) and sinusoid (see Fig. 5d). As well as the spatial frequency of a fingerprint varying naturally, the spatial frequency of latent prints present an additional issue as they may vary due to the curvature of the surface they are present on (since the surface's distance from the camera varies). Thus, we adaptively analyse local regions of the fingerprint using superpixel segmentation, estimating the local frequency separately in each superpixel. We estimate the number of fingerprint ridges present by aligning the ridges vertically using the orientation information computed from the Fourier transform, then we mean down the columns and compute the number of peaks.

D. Gray-level co-occurrence matrix (GLCM)

We also compute the gray-level co-occurrence matrix (GLCM) for each superpixel, which is a histogram of co-

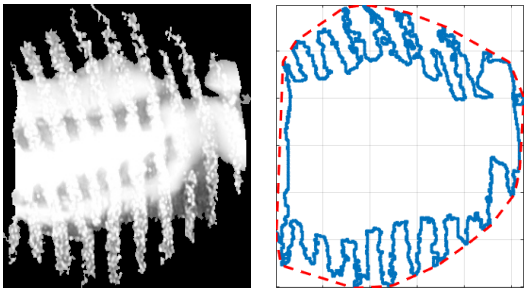

(a) Example of superpixel with meandering perimeter due to fingerprint.

(b) Perimeter and convex hull shown by solid blue and dashed red lines respectively.

Fig. 6: Fingerprint ridges present in a superpixel result in a meandering superpixel perimeter, which is larger than the superpixel convex hull. (a) Shows an example of a superpixel with fingerprint ridges (b) Shows the perimeter is larger than the convex hull due to the fingerprint ridges.

occurring grayscale values at a given offset across an image [17]. We compute the GLCM for each superpixel to quantitatively analyse their texture, allowing us to extract numeric GLCM features such as contrast which measures the intensity contrast between a pixel and its neighbour over the superpixel. The GLCM correlation measures how correlated a pixel is to its neighbour over the whole superpixel. The GLCM energy yields the sum of squared elements in the GLCM, and we also measure homogeneity which is the closeness of elements in the GLCM to the GLCM diagonal (a texture is considered coarse if most entries in the GLCM are situated down the main diagonal). The features we extract from the GLCM are used in II-E so the machine learning algorithm may learn information from them.

We will now discuss processing the superpixel to obtain more features and build $feature\ vector$ for input into a neural network.

E. Constructing a feature vector to represent a superpixel

We obtain fourteen numerical features from the superpixel to build a one dimensional feature vector. We may use this 14x1 feature vector to represent the entire superpixel, meaning we only use 14 elements to learn from instead of the $40,000$ elements (pixels) a typical superpixel may consist of (given a 24 MP input image and 400 superpixels). We use these feature vectors to train a neural network so that the network may learn information from these metrics and which combination of these metrics represent a fingerprint superpixel and which combinations do not.

1) Cross correlation of superpixel with filter: As described in II-C, we obtain the cross correlation of the superpixel against a 2D sinusoidal filter.

2) Number of ridges present in superpixel: As described in II-C, we estimate the number of fingerprint ridges present in a given superpixel.

3) Ratio of non-zero to zero value pixels in superpixel: We compute the ratio of non-zero value to zero value pixels (i.e. the ratio of light to dark pixels). This feature indicates how saturated the superpixel is and hence indicates the likelihood that a fingerprint portion is present. The SLIC superpixels algorithm groups together specular reflection pixels due to their similar (saturated) intensities. In these instances often the entire superpixel is $whited$ out, yielding a ratio of non-zero to zero value pixels of near 1, we can safely assume that the superpixel contains little to no information. Conversely if this ratio is near 0 we may assume that the superpixel is in fact too dark to obtain information from.

4) Aspect ratio of superpixel dimensions: The aspect ratio of the superpixels dimensions is computed as this can serve as a useful indicator about the contents of the superpixel. This is because the shape of specular reflections on curved surfaces are often elongated and their dimensions are highly dissimilar. We simply calculate the superpixel height and width then take the smallest of these two dimensions and divide it by the largest, meaning that the aspect ratio is rotation invariant.

5) Ratio of perimeter over area of superpixel: The ratio of superpixel perimeter to area is also a helpful numerical feature since it indicates how the superpixel has adhered to object boundaries in the image, with largely empty superpixels having smaller (more circular) perimeters and superpixels containing fingerprint portions have a larger (more meandering) perimeter as shown in Fig. 6b.

6) Ratio of convex hull over perimeter of superpixel: The convex hull of a set of points on a plane is the smallest possible convex polygon which contains all of the points in the set. The convex hull may occasionally be equal to the perimeter of the superpixel when there is less texture in the superpixel. However, the perimeter is usually larger than the convex hull in instances where the superpixel contains fingerprint portions as is shown in Fig. 6b.

7) Variance in intensity across superpixel: We compute the variance to measure how far the set of intensity values in the superpixel deviate from their average value.

8) Median intensity value of superpixel: We compute the median since outliers do not affect this feature as much as they affect the mean, which is useful when comparing superpixels that may contain a few bright specular pixels but are overall darker.

9) Mode intensity value of superpixel: We compute the modal value of the superpixel since it is also not as affected by outliers as the mean, which is useful when a small number of bright specular pixels occur in an overall dark superpixel.

10) Entropy of the superpixel: The entropy value tells us the randomness of intensity in the distribution of the superpixel. It also provides us measure of information content, estimating the amount of information present in a superpixel.

11) Contrast (from GLCM): As described in II-D, we obtain a contrast value from the gray-level co-occurrence matrix (GLCM) which is a statistical method for examining texture.

12) Correlation (from GLCM): The gray-level co-occurrence matrix (GLCM) correlation measures how

correlated a pixel is to its neighbour over the whole superpixel (see II-D).

13) Energy (from GLCM): The energy value of the gray-level co-occurrence matrix (GLCM) yields the sum of squared elements in the GLCM (see II-D).

14) Homogeneity (from GLCM): Homogeneity is a measure of the closeness of elements in the gray-level co-occurrence matrix (GLCM) to its diagonal (a texture is considered coarse if most entries in the GLCM are situated down the main diagonal). For more detail see II-D. We can now use this 14x1 feature vector to represent an entire superpixel.

F. Artificial neural network

We train a two-layer feed-forward artificial neural network to learn from 251,154 superpixel feature vectors, of which 0.6% are classified as fingerprints. These superpixel feature vectors are extracted from 13 different image stacks (with each stack containing of the order of 80 images) of varying surface specularity and colour. None of the image stacks used to generate fingerprint images in III are used during training. We use 70% (175,808) of these vectors to train the network using Bayesian regularisation [18], [19] to update the weights and biases according to Levenberg-Marquardt optimisation, minimising squared errors and weights resulting in a network with good generalisation (its ability to handle unseen data). We use 15% (37,673) of the feature vectors for validation in order to measure the generalisation of the network, and stop training when generalisation ceases to improve. We also use 15% (37,673) for testing, allowing us to evaluate the network's performance independently of the training and validation data. The feed-forward network uses a hyperbolic tangent sigmoid transfer function in its hidden layer (which consists of 10 hidden neurons) and a linear transfer function in its output layer. We train our neural network classifier on positive fingerprint superpixels as well as marginal superpixels that contain fingerprint portions that are less legible.

III. Results

We demonstrate our RELF method on prints from five objects of varying surface characteristics, showing that it is capable of extracting fingerprints from planar, curved and spherical specular surfaces. Our network produced a receiver operating characteristic curve with an area under curve of 0.9557, showing that it distinguishes well between fingerprint and non-fingerprint superpixels. Some of the results are hand-lit and some are dome-lit, in order to show RELF performs well with either. As noted in I-A, when the principal curvatures κ_1 and κ_2 increase from zero on a planar surface, to one becoming non-zero on a cylindrical surface until both are non-zero on a spherical surface, the mean curvature also increases. The mean curvature (the average of κ_1 and κ_2) of a surface is indicative of the ease at which we may extract fingerprints. We see this is true for the (hand-lit) window in Fig. 7(a)-(c), where the quality and completeness of the latent fingerprint is high. This is despite the window in question being chosen because it was particularly unclean, allowing us to evaluate the performance of RELF in a real world environment. We further evaluate RELF on a (dome-lit) specular cylindrical glass jar in Fig. 7(d)-(f). It is clear that the fingerprint is missing some portions, yet overall we yield a high quality, mostly complete latent print. This shows the robustness of RELF to work on different surface curvatures. We then evaluate RELF's performance on a clear (dome-lit) specular spherical lightbulb, as can be seen in Fig. 7(g)-(i). Since both principal curvatures are now non-zero for this surface, each input image contains only small portions of the fingerprint at best. This makes it much more difficult for RELF to extract latent prints, but the technique produces a mostly complete fingerprint nonetheless. This surface is additionally difficult to extract fingerprints as the curved and transparent lightbulb contains filament elements which are visible and disruptive to the extraction process. Despite this, RELF proves to be robust in outputting fingerprints in these most undesirable of circumstances. We also extract fingerprints from a (hand-lit) mug which has a particularly challenging combination of properties: it is black in colour, specular and curved. The fingerprint outputted from RELF is shown in 8c, which shows the stark contrast between the region of the unprocessed image in 8b. We show our novel technique is also capable of extracting fingerprints from further problematic surfaces such as the (dome-lit) white spherical specular bulb in Fig. 8f.

A Metropolitan Police Service officer has confirmed that the resulting fingerprint images are of a quality that would enable them to be matched [20]. We also evaluated our results using minutiae matching software developed by the National Institute of Standards and Technology (NIST) and used by the Federal Bureau of Investigation (FBI) [21]. The NIST biometric image software (NBIS) clusters all linked pairs of matched minutiae, and uses the length of the longest chain to provide a score. A match score of greater than 40 is stated to usually indicate a true match. The RELF output image from the flat window (7c) scored 95, the cylindrical glass jar (7f) scored 65, the spherical glass bulb (7i) scored 38 and the curved black mug (8c) scored 36 and the white bulb (8f) scored 25. These scores are promising but show that the model could perhaps be made more robust to different coloured surfaces by adapting the method in future work.

IV. Conclusions and future work

We have shown the effectiveness of the RELF technique at non-invasively extracting fingerprints from different surface curvatures, specularity and opacities without the need for chemical processing which may degrade or contaminate the evidence, preventing additional forensic testing. If a human being sat down and inspected all input images they could spot the fingerprint portion - RELF simply performs this automatically, and efficiently adds all fingerprint portions into one image uncovering the true latent fingerprint. As mentioned in II-F, we train our network on positive fingerprint superpixels as well as marginal superpixels that contain fingerprint portions which are less eligible. This produced a more complete fingerprint than only training on positive superpixels since the positive

(a) Input image of (specular) flat surface.

(b) Cropped region of (a).

(c) RELF data overlayed on cropped region of (a).

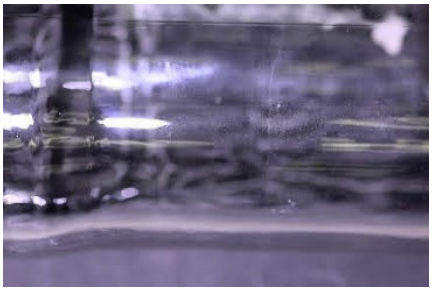

(d) Input image of (specular) clyindrical surface.

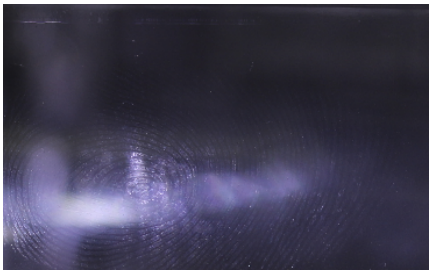

(e) Cropped region of (d).

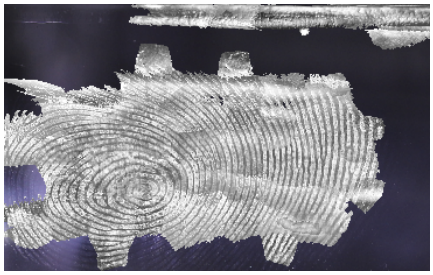

(f) RELF data overlayed on cropped region of (d).

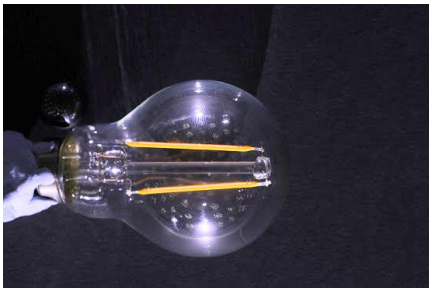

(g) Input image of (specular) spherical surface.

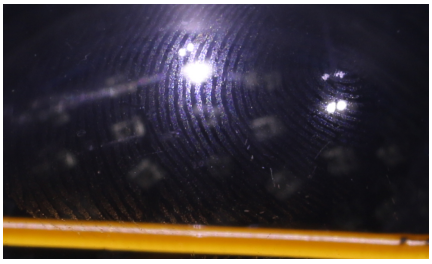

(h) Cropped region of (g).

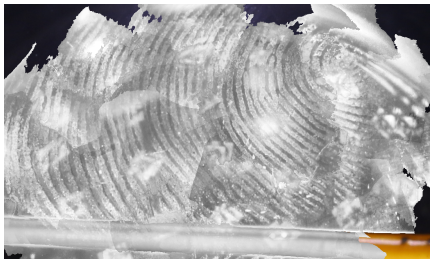

(i) RELF data overlayed on cropped region of (g).

Fig. 7: Fingerprints extracted using the RELF method on planar, cylindrical and spherical specular surfaces. It can be seen that there are portions of the fingerprint missing, due to the mis-classification of a small number of superpixels for which lighting may not have been adequate or the SLIC superpixels algorithm may have segmented a small number of unusually shaped superpixels. In (a) we see the flat surface of a window looking out with some greenery in the background. (b) shows a cropped region of (a) for comparison. (c) shows the cropped region in (b) with RELF data overlayed. In (d) we see the cylindrical surface of a glass jar. (e) shows a cropped region of (d) for comparison. (f) shows the cropped region in (e) with RELF data overlayed. In (g) we see the spherical surface of a glass lightbulb. (h) shows a cropped region of (g) for comparison. (i) shows the cropped region in (h) with RELF data overlayed. (a),(d) and (g) have been brightened, for clarity.

superpixels are very bright and result in the network becoming more dependant on simply how bright the superpixel is. We also note that we have not performed feature selection to identify features which contribute most to the prediction of a fingerprint and on which surfaces, so there may be room for optimisation to remove less relevant features to increase accuracy. There is also room for adapting the model so the method produces higher matching scores in minutiae matching software [21] across different surfaces.

We also note that the effect outlined in II-C where the curvature of fingerprint ridges increase towards the centre of the fingerprint [16] has been largely negated by using a sufficiently high number of superpixels. However, some very small central fingerprint portions are in fact missing in some of the final fingerprint images shown in 7i when both principal curvatures are non-zero making the surface more difficult to image.

As stated in I, multi-light imaging is stated to run into difficulty when imaging specular surfaces. We show that multi-light imaging is capable of extracting fingerprints from planar, curved and spherical specular surfaces. Moreover, we show the method is capable of extracting fingerprints from these difficult specular surfaces when they are transparent and unclean (see Fig. 7c), black (see Fig. 8c) and white (see Fig. 8f).

We present a method which is simple, fast and requires only an off-the-shelf camera and remote flash. Our technique is performed in real world environments with little or no preparation such as the unclean window shown in Fig. 7c,

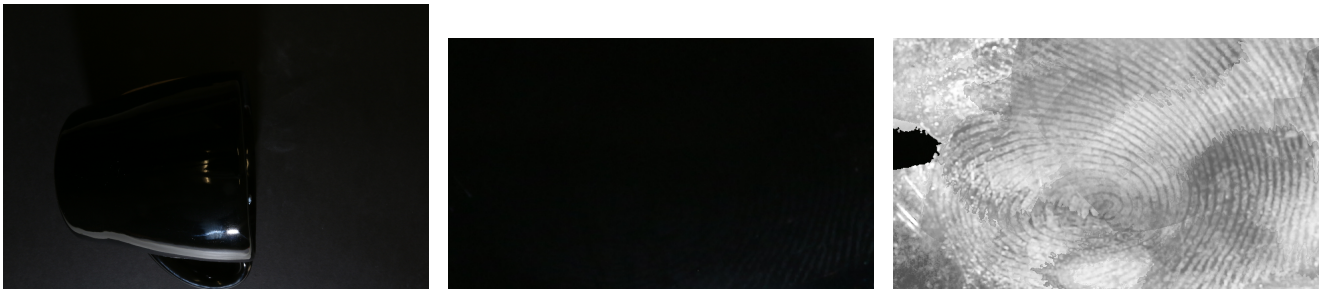

(a) Input image: black, specular, curved. (b) Cropped region of (a). (c) RELF data overlayed on cropped region of (a).

(d) Input image: white, specular, spherical (e) Cropped region of (a). (f) RELF data overlayed on cropped region of (a).

Fig. 8: Fingerprints extracted using the RELF method on a black curved specular surface (a mug) and white specular spherical bulb. In (a) we see the specular curved black surface of a mug. (b) shows a cropped region of (a) for comparison. (c) shows the cropped region in (b) with RELF data overlayed. In (d) we see the white spherical surface of a glass lightbulb. (b) shows a cropped region of (d) for comparison. (f) shows the cropped region in (e) with RELF data overlayed.

opening up the technique to potentially be used in crime scenes. We are undertaking a project with the Metropolitan Police Service which will lead to further quantitative analysis of our results and explore the method's limits with different surface types, as well as developing the mathematics behind the fingerprint classification [20].

REFERENCES

[1] L. O'Gorman, "An overview of fingerprint verification technologies," *Information Security Technical Report*, vol. 3, no. 1, pp. 21–32, 1998.

[2] M. U. Munir and M. Y. Javed, "Fingerprint matching using ridge patterns," in *2005 International Conference on Information and Communication Technologies*. IEEE, 2005, pp. 116–120.

[3] S.-S. Lin, K. M. Yemelyanov, E. N. Pugh Jr, and N. Engheta, "Polarization-based and specular-reflection-based noncontact latent fingerprint imaging and lifting," *JOSA A*, vol. 23, no. 9, pp. 2137–2153, 2006.

[4] Q. Zheng, A. Kumar, and G. Pan, "Contactless 3d fingerprint identification without 3d reconstruction," in *2018 International Workshop on Biometrics and Forensics (IWBF)*. IEEE, 2018, pp. 1–6.

[5] Z. Wang, S. Chen, C. Busch, and X. Niu, "Performance evaluation of fingerprint enhancement algorithms," in *2008 Congress on Image and Signal Processing*, vol. 3. IEEE, 2008, pp. 389–393.

[6] S. K. Dubey, T. Anna, C. Shakher, and D. S. Mehta, "Fingerprint detection using full-field swept-source optical coherence tomography," *Applied Physics Letters*, vol. 91, no. 18, p. 181106, 2007.

[7] S. Chang, Y. Cheng, K. V. Larin, Y. Mao, S. Sherif, and C. Flueraru, "Optical coherence tomography used for security and fingerprint-sensing applications," *IET Image Processing*, vol. 2, no. 1, pp. 48–58, 2008.

[8] K. Kuivalainen, K.-E. Peiponen, and K. Myller, "Application of a diffractive element-based sensor for detection of latent fingerprints from a curved smooth surface," *Measurement Science and Technology*, vol. 20, no. 7, p. 077002, 2009.

[9] L. Yan and J. Chen, "Non-intrusive fingerprints extraction from hyperspectral imagery," in *2018 26th European Signal Processing Conference (EUSIPCO)*. IEEE, 2018, pp. 1432–1436.

[10] W. R. Scott, *Fingerprint mechanics, a handbook: fingerprints from crime scene to courtroom*. Thomas, 1951.

[11] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial texture maps," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 2001, pp. 519–528.

[12] A. Goy, K. Arthur, S. Li, and G. Barbastathis, "Low photon count phase retrieval using deep learning," *Physical review letters*, vol. 121, no. 24, p. 243902, 2018.

[13] T. Tuytelaars, K. Mikolajczyk *et al.*, "Local invariant feature detectors: a survey," *Foundations and trends® in computer graphics and vision*, vol. 3, no. 3, pp. 177–280, 2008.

[14] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "SLIC superpixels compared to state-of-the-art superpixel methods," *IEEE transactions on pattern analysis and machine intelligence*, vol. 34, no. 11, pp. 2274–2282, 2012.

[15] R. Achanta, A. Shaji, K. Smith, A. Lucchi, Pascal, , and S. Süsstrunk, "SLIC superpixels," École Polytechnique édrale de Lausanne (EPFL), Tech. Rep., 2010.

[16] B. G. Sherlock, D. Monro, and K. Millard, "Fingerprint enhancement by directional fourier filtering," *IEE Proceedings-Vision, Image and Signal Processing*, vol. 141, no. 2, pp. 87–94, 1994.

[17] R. M. Haralick, K. Shanmugam, and I. H. Dinstein, "Textural features for image classification," *IEEE Transactions on systems, man, and cybernetics*, no. 6, pp. 610–621, 1973.

[18] D. J. MacKay, "Bayesian interpolation," *Neural computation*, vol. 4, no. 3, pp. 415–447, 1992.

[19] F. D. Foresee and M. T. Hagan, "Gauss-newton approximation to Bayesian learning," in *Proceedings of International Conference on Neural Networks (ICNN'97)*, vol. 3. IEEE, 1997, pp. 1930–1935.

[20] J. Cox, Private correspondence, Dec. 2019.

[21] *NIST Biometric Image Software (NBIS)*, [Accessed: 8-May- 2020]. [Online]. Available: https://www.nist.gov/services-resources/software/nist-biometric-image-software-nbis