

Laser Speckle Photography for Surface Tampering Detection

YiChang Shih, Abe Davis, Samuel W. Hasinoff^{*}, Frédo Durand, William T. Freeman
MIT CSAIL

Abstract

It is often desirable to detect whether a surface has been touched, even when the changes made to that surface are too subtle to see in a pair of before and after images. To address this challenge, we introduce a new imaging technique that combines computational photography and laser speckle imaging. Without requiring controlled laboratory conditions, our method is able to detect surface changes that would be indistinguishable in regular photographs. It is also mobile and does not need to be present at the time of contact with the surface, making it well suited for applications where the surface of interest cannot be constantly monitored.

Our approach takes advantage of the fact that tiny surface deformations cause phase changes in reflected coherent light which alter the speckle pattern visible under laser illumination. We take before and after images of the surface under laser light and can detect subtle contact by correlating the speckle patterns in these images. A key challenge we address is that speckle imaging is very sensitive to the location of the camera, so removing and reintroducing the camera requires high-accuracy viewpoint alignment. To this end, we use a combination of computational rephotography and correlation analysis of the speckle pattern as a function of camera translation. Our technique provides a reliable way of detecting subtle surface contact at a level that was previously only possible under laboratory conditions. With our system, the detection of these subtle surface changes can now be brought into the wild.

1. Introduction

Many scenarios, including law enforcement and security, require detecting whether physical objects have been tampered with, *e.g.*, [5]. Often, the contact is subtle and cannot be detected with the naked eye or by comparing pairs of before and after photographs (*e.g.*, Fig. 1(d) and (e)). We propose a new technique to detect surface changes for cases where traditional imaging is insufficient. We make use of

the speckle generated by laser illumination and exploit the fact that the precise speckle pattern observed from a given viewpoint depends on the phase of the light wavefront and, therefore, is sensitive to tiny perturbations of the imaged surface (Fig. 1(a-c)).

We focus on the situation where surface tampering is subtle, where only the phase, and not the intensity of the light reaching the camera might be altered. To address this problem, we leverage laser speckle imaging (Fig. 1). A laser speckle image encodes phase information, because speckle originates from the constructive and destructive interferences of waves reflected at different points of the surface (Fig. 2(a)) [12]. Phase differences come from the variation in travel distance, which is affected by tiny changes in the surface geometry. If the surface profile is altered by an amount as small as the laser wavelength, the speckle pattern is modified (Fig. 2(b)).

Our work is inspired by the use of speckle patterns for the measurement of fluid velocity [17, 9, 4], transparent object movement [21], motion sensing [22], and paper authentication [5]. Most prior work, however, deals with displacement parallel to the image plane and requires rigid and controlled laboratory settings. In contrast, we seek to detect *out-of-plane* modifications to the geometry that arise naturally from surface contact. Moreover, we want the ability to take a reference image of the surface, remove the imaging setup from the scene, and return later to take a new image that will reveal whether the surface was touched or not.

This latter requirement is especially difficult because speckle patterns are extremely sensitive to the position of the camera. As our experiments show, before and after photographs of a surface typically need to be taken from within about half a millimeter for our verification method to succeed. While this tolerance is straightforward to achieve in laboratory settings, using carefully calibrated static cameras, for most practical applications of surface verification, the camera and laser often cannot remain fixed in front of the surface. To achieve high-precision viewpoint alignment, we use a combination of computational rephotography, a technique to guide camera to the desired viewpoint based on relative pose estimation [2], and a new analysis of speckle correlation.

^{*}Current affiliation: Google Inc. Part of the work was also done while S. Hasinoff was with TTIC.

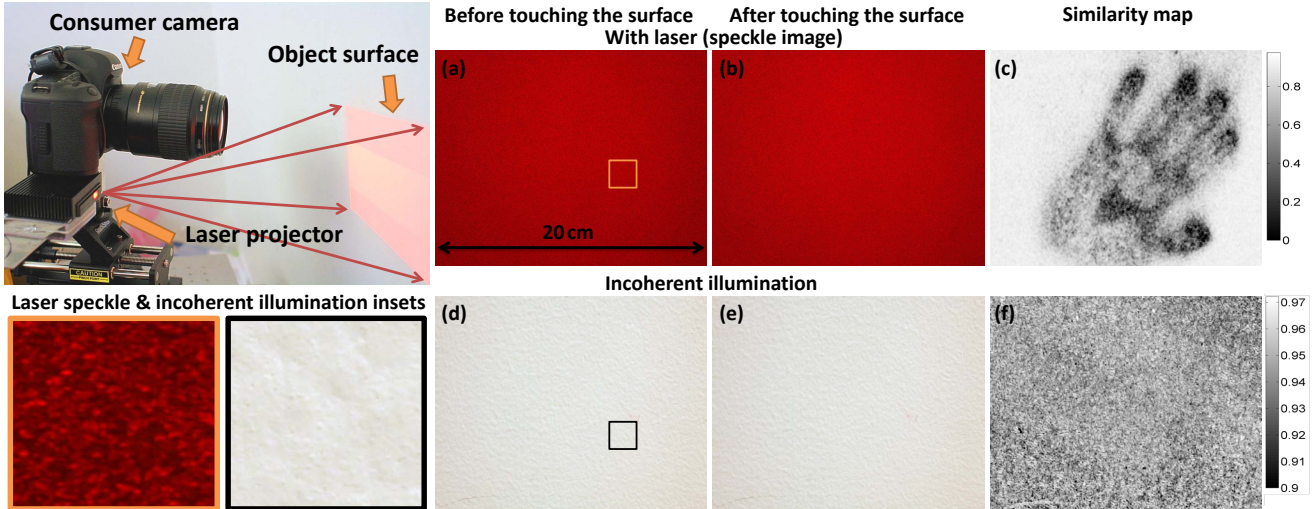


Figure 1: We detect subtle surface changes that cannot be seen in traditional photography. Top left: our proposed prototype combines an SLR with a consumer pico laser projector. (a),(b) Images of a wall illuminated by the laser projector. The granular pattern (bottom left), called speckle, is caused by interference patterns of reflected coherent light. Between (a) and (b), the wall was touched gently. The speckle similarity map we compute in (c) reveals where the wall was touched. (d)–(f): Without the laser projector, the before and after images (d) and (e) reveal no difference, as shown in the similarity map (f).

Our paper makes the following contributions:

- We present a speckle imaging system that is portable and can be moved in practice. After taking a *reference* image, the user can remove all equipment from the scene, and return later to take a *test* image. We display a visualization allowing the user to decide whether the surface has been tampered with.
- We present a new method achieving viewpoint alignment within 0.5 mm using a combination of computational rephotography and speckle correlation analysis.




Figure 2: Touching the surface causes deformation of its microstructure, which affects the wavefront of the scattered laser. The surface height before touching the surface (a) is different than afterwards (b). This affects the relative position of x_1 , x_2 , x_3 , and changes the resulting speckle image.

1.1. Related Work

Paper authentication The methods closest to our technique are those developed for speckle-based paper authentication [16, 19, 5, 20, 6, 18], in which the roughness pattern for individual pieces of paper, and their corresponding speckle, are used as an identifying signature. In one representative method, the paper must be permanently fixed to an elaborate mechanical plate that fits in a magnetic mount, to ensure accurate viewpoint reproduction [18] (p. 92). Other methods require the object to be moved and placed against a scanner. In contrast, we achieve viewpoint alignment using re-photography and speckle correlation, which alleviates the need for contact or mechanical mounts, and allows us to even handle large, immovable objects, such as a wall, the floor, and a statue.

Out-of-plane deformation Speckle interferometry can be used to measure out-of-plane deformation (*i.e.*, along the camera optical axis), but current methods require a reference laser under highly calibrated settings, for example, using a beam splitter [7]. We also seek to detect out-of-plane deformations, but with a simpler setup. Our approach can forgo the reference laser because detecting tampering does not require precisely measuring the *amount* of deformation.

In-plane motion sensing Speckle patterns can enable the use of traditional computer vision techniques to track objects such as white walls, transparent surfaces, or fluids that would otherwise be featureless. For very small translations




Figure 3: Speckle images vary over different viewpoints. The bottom row shows the image captured at each viewpoint.

parallel to the image plane, the speckle pattern is simply translated [21]. The applications include in-plane deformation measurement [15], flow visualization [17, 9, 4], and motion sensing [22]. In contrast, our work deals with out-of-plane displacement and camera positioning.

Phase retrieval Since laser speckle depends on the geometry at the scale of the laser wavelength (for us, around $1\mu\text{m}$) it may be possible to infer surface geometry from the observed pattern. This inverse problem has been explored, but the phase retrieval this requires is still an open problem [10, 13, 3, 8].

Forensics Skin contact usually leaves traces of body fluids that can be detected using fluorescent fingerprint powder and UV light, or other chemical methods. These approaches are limited to bare-skin contact and non-porous surfaces.

1.2. Overview

Our main imaging setup relies on a standard digital SLR and a consumer micro laser projector (Fig. 1). We currently use a second camera for rephotography, mounted rigidly to the imaging system (Fig. 12). In principle, a single camera could be used for both rephotography and speckle imaging, but we found that using a second camera helped us sidestep low-level software issues. The two cameras and projector are rigidly mounted on actuators to enable precise displacements during the final phase of viewpoint alignment. In our prototype, the setup is placed on a cart for easier portability.

The user first takes a reference speckle image where the surface is illuminated by the laser projector. The user then leaves with the camera-projector setup. The surface might then be tampered with or not, and the goal is to determine whether tampering has occurred. Later, the user comes back and needs to take a new speckle image from the same location to detect if the surface was touched. We first use vision-based rephotography [2, 1] to guide the camera to a

roughly accurate viewpoint within 3 mm. For flat surfaces where vision-based pose estimation is degenerate, we use structured light and the visual alignment of a checkerboard pattern instead. In both cases, we then perform a novel correlation analysis between the current laser speckle image and the reference. The correlation depends on the distance between the viewpoints and lets us provide feedback to the user to help them move even closer to the desired viewpoint.

We display the results of tampering detection by running normalized cross correlation between the reference and the final test speckle image (Fig. 1c).

2. Speckle Image Formation and its Variability

We review the equations of speckle image formation [11, 12] in order to analyze how the speckle pattern varies with camera translation and aperture choice. Detailed derivations are provided in the supplementary material. The one dimensional speckle image $I(y)$ caused by a surface with surface height $h(x)$ is

$$I(y) = \left\| f\left(-\frac{z_1}{z_2}y\right) \otimes \tilde{g}\left(\frac{y}{z_2\lambda}\right) \right\|^2, \quad (1)$$

where

$$f(x) = \frac{A_0}{z_1} e^{jk\left(h(x) + \frac{x^2}{2z_1}\right)} \quad (2)$$

is the input function, x and y are the coordinates on the object and image plane respectively, A_0 is the incident wave amplitude, z_1 and z_2 are distances from the lens to the object and to the sensor respectively, and λ is the incident wavelength. We use \otimes to denote convolution and $\tilde{g}(\omega) = \mathbb{F}\{P\}$ is the Fourier transform of the aperture function $P(u)$.

Eq. (1) decomposes speckle formation into three steps: (1) the input $\frac{A_0}{z_1} e^{jk h(x)}$ is modulated by a quadratic phase $e^{jk \frac{x^2}{2z_1}}$ and an attenuation $\frac{1}{z_1}$; (2) it passes through a low-pass filter $\tilde{g}(\omega)$ determined by the aperture; (3) the speckle image intensity is the squared amplitude of the field. The low-pass filter (2) comes from the fact that the field at the lens is the Fourier transform of the surface field and it gets clipped by the aperture. This model demonstrates locality: the modification of the surface $h(x)$ at a point x only affects the image near $-\frac{z_2}{z_1}x$ because a finite aperture causes \tilde{g} to fall off rapidly, which is important for our application.

2.1. Speckle variation with camera movement

Camera in-plane translation When the camera moves by δ in the lens plane, we simply need to replace the aperture in Eq. (1) by $P(u - \delta)$. As the camera moves, the aperture selects different windows over the field. If the translation is larger than the aperture, the windows don't overlap




Figure 4: Laser speckle captured with larger apertures has more high frequency components but lower contrast.

(Fig. 3). Furthermore, these windows are generally uncorrelated because, in practical applications, the surface height $h(x)$ is randomly distributed, and the corresponding speckle images are also uncorrelated. For our surface tampering detection, we need the speckle from untouched regions to be highly correlated in the before/after pair, and the captured windows of the field must be very close.

To understand the effect of camera translation, we analyze the case where the camera remains static and the surface is translated by $-\delta$. The input f in Eq. (1) is translated, but the quadratic phase term $\frac{x^2}{2z_1}$ in $f(x)$ in Eq. (1) makes the system not shift-invariant and the speckle is altered.

Depth translation With the camera in focus, Eq. (1) shows that changing the lens-object distance z_1 causes speckle magnification and alteration. Fortunately, this alteration is not as strong as that associated with in-plane translation. The quadratic phase term $e^{jk\frac{z_1 y^2}{2z_1^2}}$ is less affected by depth z_1 than by in-plane translation y . Hence, the required precision for viewpoint alignment is looser along depth and can be achieved with vision-based rephotography alone.

Aperture size Speckle images taken at different apertures can be derived by evaluating Eq. (1) with appropriate $P(u)$. From Eq. (1), the aperture is equivalent to a low-pass filter, and larger apertures correspond to filters with larger bandwidth. This is equivalent to the well-known reduced diffraction with large apertures. As shown in Fig. 4, images with smaller apertures have lower frequency speckle.

However, increasing aperture size reduces speckle contrast, which leads to a tradeoff, since contrast is needed to identify tampering. In our work, we use a 100 mm lens and found that apertures between $f/6.7$ and $f/16$ work well.

3. Rephotography and Similarity Map Computation

As shown in Fig. 5, the reference and test images become decorrelated when the viewpoints differ by more than

1 mm, which sets tight accuracy goals for camera alignment. To our knowledge, no computational rephotography [2] technique can achieve such accuracy. We propose a two-stage method: first, we adopt vision-based rephotography to get within a few mm, and second, we introduce a speckle-based correlation analysis for finer accuracy.

3.1. Vision-based computational rephotography

For the first stage, we use two rephotography techniques for different types of scenes. For general 3D scenes with distinctive image features, we use a solution similar to Bae et al. [2] and rely on these features to estimate the relative camera pose. For flat scenes, where pose estimation is degenerate, we project structured light onto the surface and rely on the user to visually match the observed projected pattern. We found that both of these techniques have similar accuracy, and can typically guide the camera back to the reference viewpoint to within a translation of 3 mm and a rotation of 0.5 degrees (Table 1).

Image-based feature matching and pose estimation

Our feature-based approach builds on previous methods for feature detection and camera pose estimation, namely the parallel tracking and mapping (PTAM) library [14]. PTAM builds a 3D map where feature points are characterized by 2D image patches. It uses a combination of tracking, RANSAC and bundle adjustment to continuously refine the map and compute the pose of the current camera.

Prior to taking the reference image, we use PTAM to build up a sparse 3D map of feature points from a set of uncalibrated wide-angle photos surrounding the object. We then record the camera pose, turn on the laser projector, and take the reference speckle image.

When we return to take the test image, the current camera pose is estimated by PTAM using the earlier map. We display a 3D visualization of the current displacement from the reference pose, and manually adjust the camera viewpoint until the pose returns to the reference viewpoint. This method works well when the scene has enough features and depth range. It typically takes us a few minutes to reach the viewpoint with 3 mm accuracy.

Structured light for flat scenes

Flat scenes, such as walls, present a special challenge for viewpoint alignment since small camera rotations and translations are ambiguous for flat surfaces. This ambiguity limits viewpoint alignment, which is why we propose to project structured light onto the surface to help resolve the ambiguity. In practice, we use the laser projector to project a checkerboard pattern onto the scene and capture an image of this pattern from the reference viewpoint. At test time, matching the appearance of the projected checkerboard lets us resolve the translation-rotation ambiguity. Because the camera and the projector




Figure 5: Top: Sensitivity of speckle correlation to viewpoint change, as measured using our system under different aperture sizes.

are not co-located, the resulting parallax deforms the observed pattern when the system is moved. However, when the system translates in a direction parallel to the flat surface, the observed pattern is invariant. In order to resolve these remaining two degrees of freedom, we use a single 2D feature on the surface, *e.g.*, a door knob. We use the same projector for speckle imaging so only one is needed for the whole system. Note that the geometric relation between the camera and the projector is fixed, so there is no transform between taking the reference pattern and taking the test pattern. The alignment is currently based on visual matching and manual displacement.

3.2. Speckle-based viewpoint alignment

In the second stage of viewpoint alignment, we focus on refining the in-plane camera translation by analyzing the correlation between the current speckle image and the reference one. This provides users with feedback as to whether they are getting closer or farther from the reference viewpoint. As described in Section 2.1, speckle correlation is less sensitive to depth change, so the first stage of viewpoint alignment provides sufficient accuracy for this degree of freedom.

Within the uncertainty range given by the first stage, we sample along the camera plane (perpendicular to the optical axis). For each sampled image we display its normalized cross correlation (NCC) with the reference speckle image as feedback to the user. As shown in Fig. 5, NCC reaches its maximum when the viewpoints are the same. Hence, by maximizing the NCC, we can determine the desired position. In our work we sample every 0.5 mm in the range [-5 mm, 5 mm]. When computing the NCC, we first crop a patch in the sampled image and then search for the patch in the reference image that results in the maximum NCC with our sampled patch. The NCC value is displayed as feedback to the user. This helps the user resolve any misalignment between two speckle images due to the small rotation uncertainty given by the first stage.

After determining the camera translation, the user corrects the camera rotation by matching 2D features from the

experiment		translation, horizontal (mm)	translation, vertical (mm)	rotation (arcmin)
Statue	(F)	2.5	0.2	23
Cup	(F)	2.1	0.3	20
Cylinder	(F)	1.3	0.5	17
Door	(S)	1.2	0.1	12
Wall	(S)	1.1	0.2	13
Drawer	(S)	1.7	0.1	18

Table 1: Accuracy of vision-based rephotography, measured by the amount of viewpoint change required by the second stage of our viewpoint alignment. Experiments marked (F) used the feature-based PTAM method, while those marked (S) used our structured light technique.

current frame with the reference frame.

3.3. Similarity Map Computation

Once the correct viewpoint has been reached, we take the test speckle image. We carefully align this image to our reference image before analyzing their differences.

To match the speckle images we randomly sample image patches, match those patches, and fit a warp to the resulting patch correspondences. As Eq. (1) shows, touching the surface only affects the speckle image near the touched area, so local comparisons are sufficient for matching. First, we sample several hundred patches from the reference image, then for each patch we use normalized cross correlation to find the corresponding best-matching patch in the test image. Given these correspondences, we use RANSAC to recover a 3×3 homography relating the two speckle images, while ignoring potential outliers due to tampering.

We then compute the similarity map S relating the target image I_{tar} and the warped reference image $I_{\text{ref,w}}$,

$$S(i, j) = NCC(W(i, j)I_{\text{tar}}, W(i, j)I_{\text{ref,w}}), \quad (3)$$

where NCC computes the normalized cross correlation, and $W(i, j)$ is a windowing function centered at (i, j) . For our results we used a window size of 21×21 pixels.

4. Hardware

Our prototype is shown in Fig. 6. We use the Microvision PicoP® as a laser projector and the Canon EOS 5D Mark II® for capturing the speckle image. For camera registration, we use a motorized labjack to adjust the vertical position of the camera, two motorized translation stages to adjust the depth and horizontal position, and a motorized rotation stage to adjust camera rotation. The precision of the labjack and translation stage is 0.01 mm, and the rotation stage has a measurement precision of 1 degree but a higher precision for actuation. We put the whole system on a cart for mobility.




Figure 6: System prototype. Our system consists of a controller with 4 degrees of freedom. The SLR and the laser projector are fixed on the controller.

5. Results

We performed a number of tests to evaluate surface tampering detection and viewpoint alignment using our system.

First, we tested the sensitivity of our similarity map (Eq. (3)) to demonstrate its ability to detect subtle surface changes resulting from small forces. To study this sensitivity independent of issues with viewpoint alignment, we kept the camera fixed on a tripod while photographing the reference (before surface deformation) and test (after) images. In the experiment of Fig. 7a, we placed a quarter (5.67 g) on a heavy cardboard box. The similarity map of the test to reference images clearly reveals the effect of the quarter’s weight on the box. Only the boundary of the quarter is visible because the coin is thicker at the periphery and does not touch the surface in the center. To avoid extra forces on the box we taped the coin to a string, lowered it slowly onto the box, and removed it by lifting the string.

The experiment of Fig. 7b reveals the effect of touching a sheet rock wall with a rubber square using the lightest force we could measure with a hand-held pressure scale apparatus. The mark from the 20.7 gf/cm² pressure (gf = gram-force) is clearly visible in the normalized correlation similarity map.

Fig. 8 shows the results of our system detecting finger-pressure marks on surfaces of different shapes, colors, and materials.

The remaining figures demonstrate the reliability of our viewpoint alignment schemes for surface tampering detection. Fig. 10 shows the effectiveness of the structured light method for viewpoint alignment with flat scenes. The surface tampering detection reveals light finger touches to a wooden door, a wall, and a steel filing cabinet.

Figs. 11 and 12 show the results of using our vision-based rephotography to reveal light finger marks on 3D surfaces.

The average time for two-stage camera alignment in both cases is around 5 minutes. Typically, both of the first-stage




Figure 7: Our technique can detect tiny surface variations due to light pressure. (a) We put a US quarter dollar coin (weight 5.67 g) onto a box, then removed it. The similarity map relating the before and after speckle images reveals the outline of the coin. (b) We attach a piece of incompressible rubber to a sensitive scale. We press lightly against the wall with the rubber, and zero the scale while touching. The readout from the scale shows that the force is 27 g. Since the surface area of the rubber is 1 cm × 1.3 cm, the pressure is 20.7 gf/cm². The similarity map relating speckle images before and after touching the wall shows that the touched area can be identified.

rephotography options, feature-based PTAM and our structured light technique, allow us to bring the camera back to within 3 mm, as shown in Table 1.

Comparison with UV light and fingerprint powder We compare our work against a leading forensic technique for fingerprinting that uses fluorescent powder. When a human touches an object, some fluids from the finger are transferred to the object’s surface. Later, the surface can be brushed with a fluorescent powder that is adhesive to fluid. A Shining a UV light on the brushed surface will then reveal the pattern of fluid left by the finger.

As shown in Fig. 9, our method and the fingerprint powder technique have different limitations. Our approach does not work on transparent objects because there is no reflected laser light, whereas fingerprint powder works well. Porous surfaces, on the other hand, present the opposite scenario because fluids do not stick to them well, and, moreover, the powder tends to stick even in the absence of fluid. In contrast, our method works well for such surfaces. Fig. 9 also shows that the fingerprint powder fails to reveal the touch of a gloved hand, while the speckle change is clearly visible.

Fingerprint powder does not require a reference image and is often able to recover the detailed groove pattern of a fingerprint. However, unlike our technique which does not require any contact with the scene, fingerprint powder is potentially destructive to the surface being examined.

5.1. Limitations

We showed that our system works well for a variety of common materials. However, we have found that our system struggles with the following three types of materials. (1) Volumetric materials, such as carpets, sweaters, and scarves. For these materials, the surface is not well defined,

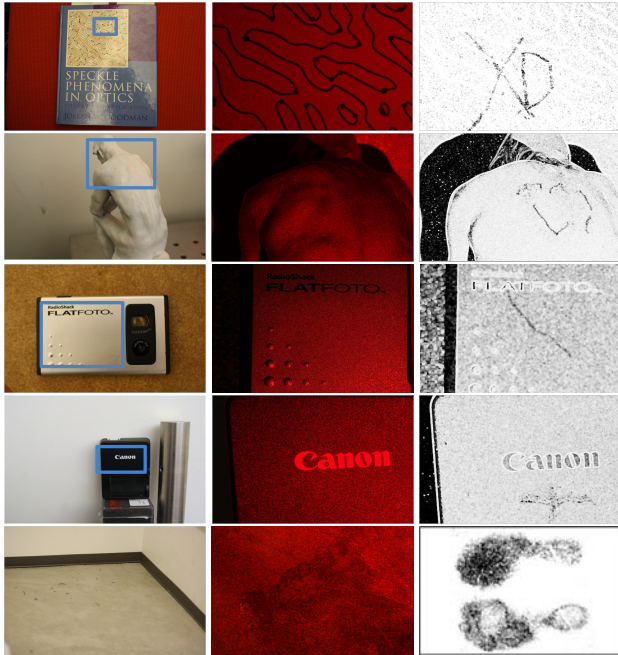


Figure 8: Surface tampering detection on different materials using our system. From top to bottom: a glossy textbook cover with a textured design (we drew an XD on it with the fleshy part of a finger), a stone statue (drew a smile), a brushed aluminum camera (drew a backslash), a black plastic camera charger (drew a cross), and a cement floor (stood on the ground wearing shoes). Speckle is still observable for black objects provided that the exposure time is long enough. Our technique also works well for walls and boxes (Figs. 7a)

and our camera cannot observe speckle. (2) Reflective and transparent materials, such as mirrors, shiny metals or glass. (3) Materials that absorb light at the bandwidth of our laser, such as cyan material for a red laser, or black materials for all bandwidths. In general, diffuse and moderately-glossy surfaces that are not too dark work best.

6. Discussion

In this work, we presented a technique that can non-invasively detect otherwise-invisible surface tampering by comparing laser speckle images. Using a combination of vision-based rephotography, structured light, and speckle correlation analysis, we can guide the user within the required viewpoint accuracy of 0.5 mm, typically within five minutes. Our method allows for small surface changes to be detected under very general conditions

Acknowledgements

We thank MicroVision for donation of equipment, and acknowledge gifts from Microsoft Research and Texas Instruments, and funding from NSF CGV No.1111415.




Figure 9: We compare our technique to a forensic fingerprinting technique that uses fluorescent powder. We brush fingerprint powder on the object to detect the latent fingerprint. Our method struggles with transparent objects, while fingerprint powder often fails on porous surfaces. Bottom 2 rows: In the top row we use a bare finger to touch the surface, while in the bottom row we touch the surface with a gloved finger. Fingerprint powder does not work in the latter case, while our technique works in both cases.

References

- [1] A. Adams, D. Jacobs, J. Dolson, M. Tico, K. Pulli, E. Talvala, B. Ajdin, D. Vaquero, H. Lensch, M. Horowitz, et al. The Frankencamera: An experimental platform for computational photography. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 29:29:1–29:12, 2010. 3
- [2] S. Bae, A. Agarwala, and F. Durand. Computational rephotography. *ACM Transactions on Graphics*, 29:24:1–24:15, 2010. 1, 3, 4
- [3] H. Bauschke, P. Combettes, and D. Luke. Phase retrieval, error reduction algorithm, and Fienup variants: a view from convex optimization. *JOSA A*, 19(7):1334–1345, 2002. 3
- [4] J. Briers and S. Webster. Laser speckle contrast analysis (LASCA): a non-scanning, full-field technique for monitoring capillary blood flow. *Journal of Biomedical Optics*, 1:174–179, 1996. 1, 3
- [5] J. Buchanan, R. Cowburn, A. Jausovec, D. Petit, P. Seem, G. Xiong, D. Atkinson, K. Fenton, D. Allwood, and M. Bryan. ‘Fingerprinting’ documents and packaging. *Nature*, 436(28):475, 2005. 1, 2




Figure 10: Camera viewpoint alignment for flat scenes: (a) door, (b) wall, (c) drawer. For each surface, we first take a reference picture, and then move the system away before gently touching the surface. We then perform our camera viewpoint alignment scheme on each flat scene to find the viewpoint for the test image. The similarity maps (a)-(c) show that the touched area can be detected by our technique.

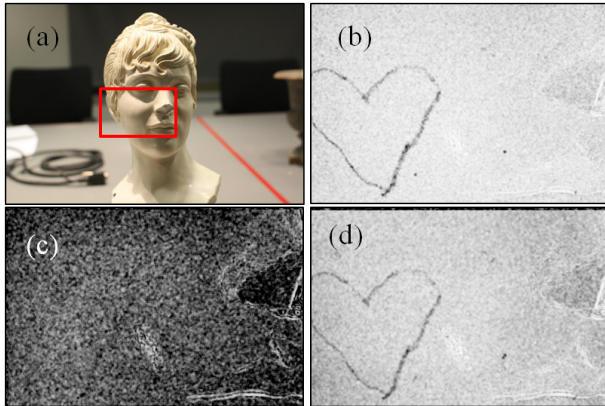


Figure 11: Camera viewpoint alignment for a statue of a woman's head (a). We take a reference image, and use a finger to draw a heart on the woman's cheek (in the red rectangle). (b) We calculate the similarity between the reference image and the test images taken without moving the camera, (c) after moving the camera away and using only vision-based rephotography for viewpoint alignment, and (d) using the whole proposed viewpoint alignment method.




Figure 12: Camera viewpoint alignment for two non-planar scenes: (a) a cup (c) a highly textured cylinder. We gently touch the surfaces inside the red rectangle and remove our system from the scene before capturing test images. (b) and (d) show the similarity maps for (a) and (c). Our technique successfully detects the touched regions.

[6] R. Cowburn. Laser surface authentication-reading nature's own security code. *Contemporary Physics*, 49(5):331–342, 2008. 2

[7] K. Creath. Phase-shifting speckle interferometry. *Appl. Opt.*, 24(18):3053–3058, 1985. 2

[8] H. Faulkner and J. Rodenburg. Movable aperture lensless transmission microscopy: a novel phase retrieval algorithm. *Physical review letters*, 93(2):23903, 2004. 3

[9] A. Fercher and J. Briers. Flow visualization by means of single-exposure speckle photography. *Optics communications*, 37(5):326–330, 1981. 1, 3

[10] J. Fienup. Phase retrieval algorithms: a comparison. *Applied Optics*, 21(15):2758–2769, 1982. 3

[11] J. Goodman. Statistical properties of laser speckle patterns. *Laser speckle and related phenomena*, pages 9–75, 1975. 3

[12] J. Goodman. *Speckle phenomena in optics: theory and applications*. Roberts & Co, 2007. 1, 3

[13] S. Gravel and V. Elser. Divide and conquer: A general approach to constraint satisfaction. *Physical Review E*, 78(3):036706, 2008. 3

[14] G. Klein and D. Murray. Parallel tracking and mapping for small AR workspaces. In *Proc. Sixth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'07)*, November 2007. 4

[15] N. Mo and J. Shelton. Laser speckle photography to measure surface displacements on cortical bone—verification of an automated analysis system. *Medical engineering & physics*, 20(8):594–601, 1998. 3

[16] R. Pappu, B. Recht, J. Taylor, and N. Gershenfeld. Physical one-way functions. *Science*, 297(5589):2026, 2002. 2

[17] H. Rabal and R. Braga. *Dynamic laser speckle and applications*. CRC, 2008. 1, 3

[18] P. Ravikanth. *Physical one-way functions*. PhD thesis, Massachusetts Institute of Technology, 2001. 2

[19] P. Seem, J. Buchanan, and R. Cowburn. Impact of surface roughness on laser surface authentication signatures under linear and rotational displacements. *Optics letters*, 34(20):3175–3177, 2009. 2

[20] A. Sharma, L. Subramanian, and E. Brewer. Paperspeckle: microscopic fingerprinting of paper. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 99–110. ACM, 2011. 2

[21] R. Sirohi. *Speckle metrology*. CRC Press, 1993. 1, 3

[22] J. Zizka, A. Olwal, and R. Raskar. Specklesense: fast, precise, low-cost and compact motion sensing using laser speckle. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 489–498. ACM, 2011. 1, 3