



# TextureSight: Texture Detection for Routine Activity Awareness with Wearable Laser Speckle Imaging

XUE WANG, University of California, Los Angeles, USA

YANG ZHANG, University of California, Los Angeles, USA



Fig. 1. Our hands engage with a wide variety of objects. These objects are often in contact with the fingers, which makes for an ideal location for sensor instrumentation. *TextureSight* implements wearable laser speckle imaging to distinguish between surface profiles (i.e., textures) for object recognition, which serves as a strong context clue for higher-order activity recognition.

Objects engaged by users' hands contain rich contextual information for their strong correlation with user activities. Tools such as toothbrushes and wipes indicate cleansing and sanitation, while mice and keyboards imply work. Much research has been endeavored to sense hand-engaged objects to supply wearables with implicit interactions or ambient computing with personal informatics. We propose *TextureSight*, a smart-ring sensor that detects hand-engaged objects by detecting their distinctive surface textures using laser speckle imaging on a ring form factor. We conducted a two-day experience sampling study to investigate the unicity and repeatability of the object-texture combinations across routine objects. We grounded our sensing with a theoretical model and simulations, powered it with state-of-the-art deep neural net techniques, and evaluated it with a user study. *TextureSight* constitutes a valuable addition to the literature for its capability to sense passive objects without emission of EMI or vibration and its elimination of lens for preserving user privacy, leading to a new, practical method for activity recognition and context-aware computing.

CCS Concepts: • Human-centered computing → Ubiquitous and mobile devices.

---

Authors' addresses: Xue Wang, University of California, Los Angeles, California, USA; Yang Zhang, University of California, Los Angeles, California, USA, yangzhang@ucla.edu.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

© 2023 Copyright held by the owner/author(s).

2474-9567/2023/12-ART184

<https://doi.org/10.1145/3631413>

Additional Key Words and Phrases: laser speckle imaging, wearable computing, object detection, activity recognition, context-aware computing, smart rings

**ACM Reference Format:**

Xue Wang and Yang Zhang. 2023. *TextureSight: Texture Detection for Routine Activity Awareness with Wearable Laser Speckle Imaging*. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 4, Article 184 (December 2023), 27 pages. <https://doi.org/10.1145/3631413>

## 1 INTRODUCTION

Sensing user activities has long been researched as a key enabler in applications such as personal informatics, mobile health, and context-aware computing. Much effort in literature has been devoted to activity-sensing techniques using a wide spectrum of device form factors. Closely related to this work are prior systems worn by users that monitor users' hand activities as reliable indicators of user activity for their strong correlations (e.g., blender use indicates food preparation). Furthermore, the task of sensing hand activities could be equivalent to sensing objects that are touched, gripped, or grasped by users' hands. We call these objects hand-engaged objects. For example, if a wearable senses that its wearer is engaged with a steering wheel, they are most likely driving. Similarly, tools like screwdrivers and saws indicate fabrication works, while cookware implies a cooking activity.

To sense hand-engaged objects, one thread of research detects their activation signals such as electromagnetic radiations [30, 59] and vibrations [29]. To sense objects that do not emit active signals (e.g., mugs, bags, steering wheels), prior work has turned to vision-based approaches for solutions [34, 35]. These systems use wrist-worn cameras pointing toward a user's fingers to have sight of hand objects to be detected. Although the set of detectable objects could be highly comprehensive, this approach potentially induces concerns from users regarding privacy for the possibility of the worn camera capturing unwanted information in the background. While obfuscation techniques have been proposed [2, 51], whether these techniques could mitigate user privacy concerns in the short or long term and improve the adoption of wearable cameras is left untested.

We propose *TextureSight*, a ring sensor that recognizes hand objects by sensing their textures using laser speckle imaging (Figure 1). Our observation is that objects are comprised of unique combinations of surface color, specularity, and roughness, all of which can be detected by laser speckle imaging. This is a promising vision-based technique commonly used in the medical domain and recently demonstrated in smart-sensing systems in HCI (e.g., [9, 47, 65, 75, 77]). It is uniquely leveraged in this research for hand-engaged object detection. As laser beams scattered by surfaces are always in focus, our device does not require a lens in front of the CCD sensor, meaning laser speckles form the only signal registered on the image while blurring out any information from the background. We chose a ring form factor, which has been increasingly popular and has seen success in many prior HCI systems (e.g., [3, 41, 57, 58, 74]). More importantly, it allows our sensor to be placed in contact with sensed objects. With our mechanical design of the smart ring prototype, noise such as environmental illumination can be shielded from our sensor while ensuring its constant distance from sensed objects to improve the signal-to-noise ratio (SNR). To further improve our sensing performance, we used a dual-laser configuration and adopted a modified state-of-the-art deep neural net inference pipeline.

*TextureSight* enables smart rings to sense hand-engaged objects, allowing contextual information to be seamlessly collected without altering their users' existing activities. Once this information is collected, it could be made readily shareable with other smart wearables and computing modalities to power their context-aware applications. We envision our sensing technique to coexist with other sensors on smart rings, making these devices more capable of assisting users in their daily tasks and delivering pertinent information based on the user's specific context, ultimately contributing to a more intuitive and efficient HCI paradigm.

Our sensing technology is based on laser Speckle Imaging, which has been used conventionally in medical applications for detecting sub-dermal signals such as blood flow velocity for skin health monitoring [36, 45, 50]. It has also been recognized in HCI for its potential in high-resolution spatial tracking in interactive systems (e.g., SpecTrans [47], SpeckleSense [77], Spedo [26], and Colux [56]). Closely related to our work are previous systems that use laser speckle imaging to identify material types. Two pioneering works that inspired this research are SpecTrans [47] and SensiCut [9]. Although both works recognize material types using laser speckle imaging, they feature very different applications from our research. We focus on laser speckle imaging in sensing objects engaged by users' hands, while prior work demonstrated explicit, transparent, and specular material detection as an input technique (i.e., SpecTrans) and for improving laser cutting safety (i.e., SensiCut) with lasers at fixed heights on laser cutting machines. However, neither of the previous works have demonstrated a compact, wearable, and versatile design that worked on recognizing various surfaces with different textures and assisting users in their daily lives. The difference in target application requires different design considerations, hardware and software implementations, and systematic modeling and evaluations, all of which constitute our novel design, given the literature.

To investigate the types of hand-engaged objects, their surface materials, and the unicity and repeatability of the combinations of objects and textures, we conducted a study with six participants, collecting two days of data from each using an experience-sampling approach. We then grounded our sensing principle by modeling laser speckles, which was validated with simulations and an experimental real-world test. Results from these investigations justified our assumption that texture is of sufficient information to identify routine objects and informed our sensor implementation, which we achieved with a bare CCD sensor chip, two low-power lasers, and a 3D-printed enclosure. Finally, we conducted a user study of 11 participants using test objects informed by the survey study. Overall, our sensor achieved an average accuracy of 88.20% in the detection of 20 objects.

## 2 RELATED WORK

In this section, we first review prior systems that leveraged laser speckle imaging in interactive systems. We then move on to systems in HCI with smart ring form factors. Finally, we review sensing approaches in object detection and activity recognition.

### 2.1 Laser Speckle Imaging

The characteristics of the micro geometry on a rough surface can be extracted when coherent and collimated light falls on rough surfaces [8] and is captured by the camera when the light is scattered back, revealing the surface information and enabling interactive applications. With static laser speckles, SensiCut [9] proposed a material-sensing platform to recognize materials which are indistinguishable to the naked eye for laser cutting according to the otherness of static laser speckles on different surfaces of a flat plate. Similarly, SpecTrans [47] presented a novel surface classification method for specular, textureless, and transparent materials based on the information conveyed by static speckles. In addition, dynamic speckles caused by surface displacement are carriers of information from another dimension, which is the object's locomotion along with time. Shih et al. [53] took advantage of speckle motions to detect subtle tampering in 2D surfaces. The characteristics of dynamic laser speckles can also be implemented in medical applications [4, 11, 14]. Furthermore, laser speckle imaging is not limited to 2D motions, it can also track surface shifts in the 3D world [26]. For instance, CoLux [56] analyzed the global scale-space speckle motion to acquire information about micro objects' compound movements. Besides motion and surface perception, it is also possible to sense minute vibration signals in the environment based on the spatial change of laser speckles (i.e., laser speckle contrast imaging) [75, 76] and the intensity change of them (i.e., interferometry) [75, 76] for human activity recognition. A long-range laser supports the feasibility of low-power and city-scale sensing, which can be used for environmental perception in extra-large spaces.

Although laser speckle imaging has been demonstrated, most prior systems are in stationary settings – deployed in the environment or affixed inside machinery. Recently, LaserShoes [65] recognized ground surface types (e.g., asphalt vs. carpet) as a strong indicator of user locations for applications such as running assistance and navigation.

Closest to our work is SpecTrans [47], which detects a unique category of surfaces, such as transparent showcase, for explicit interactions, whereas our work detects a wide variety of surfaces on everyday objects including ones with smooth, rough, hard, soft, and furry surfaces and of various colors and specularities for activity recognition. To achieve this superior sensing capability, *TextureSight* leveraged a smart ring device mechanical design that yielded high SNRs, a dual-laser configuration with selected wavelengths to reveal high-fidelity texture information, and a deep learning pipeline that is best at inferring texture information from high-dimensional data types. All of these contributed to the unique sensing advantages of *TextureSight* over literature.

## 2.2 Smart Rings

Smart rings elicit many possibilities for touch sensing, gesture recognition, and remote control due to their wearability, dexterity, and maneuverability. For example, OctaRing [32] introduced an octagon-shaped ring device enabling complex multi-touch detection implemented by pressure sensors. Likewise, TouchRing [58] also allows for subtle multi-touch input using capacitive sensing, and it is capable of recognizing swipe and tap gestures based on relative touch positions of the ring. With miscellaneous sensing principles having been employed in the smart ring design, research efforts have been invested in complex gesture recognition [16] implemented via inertial measurement units [15, 22, 69], acoustic [72], magnet [3, 41], proximity [57], thermal imaging [73], electrical sensing [27], or mixed approaches [70, 71] for remote controlling. For instance, Gheran et al. [17] implemented car control buttons from a steering wheel to the smart ring without hand shifting or eye staring while driving, providing a more efficient and safer interaction via tap, touch, and mid-air gestures. Exceptions to this trend are creative gesture designs specific to some smart rings only.

Smart rings also have the capability for personal care when omnifarious sensors are embedded [12, 37], which provides near-infinite possibilities for long-term physical and mental health monitoring. For instance, the ring can be regarded as a mobile health system sensing physiological signals to assess physical and mental conditions. Wu et al. [63] proposed a wearable ring-type pulse monitoring sensor for human pulse and temperature surveillance. Multisensor smart rings were utilized to evaluate sleep quality by recording and assessing the signal of the heart rate, skin temperature, and motions [21]. Except for the perception of the signal from the human body, the wearable smart ring can also observe hand-washing duration and detect various fluid agents for monitoring hand hygiene compliance and preventing diseases [74].

Our sensing technology further expands the sensing capabilities of smart rings, providing deeper context-based insights for interaction. We uniquely enabled texture sensing with Laser Speckle Imaging using a CCD sensor in concert with two laser diodes, enabling the detection of a wide range of objects with diverse textures and colors, yielding contextual information that could be useful in many smart ring applications. Unlike prior methods, our device shows robustness to environmental factors due to the shielded sensor configuration. Furthermore, the lens-free design ensures that only speckle information is captured, respecting privacy concerns by avoiding the capturing of sensitive information.

## 2.3 Object Recognition through Wearable Sensors

Active objects emit distinguishable signals that propagate through space or the user's body and can be detected by wearables. For example, Viband [29] senses micro-vibrations from active hand-held objects propagating through the arm, while MagnifiSense [59] and EM-Sense [30] capture electromagnetic radiation emitted from objects that

run electricity. It is also possible to sense the motion of a user’s hand engaging with manipulations of objects for fine-grained activity recognition [28]. This section focuses on wearable techniques that share a similar sensing principle with our proposed work – recognizing objects through their passive material or surface properties.

One of the most common approaches relies on computer vision. Camera approaches have been used on wearables to recognize objects by their shape, color, and coarse texture [34]. G-ID [10] demonstrated that it is possible to encode information into the surface textures of 3D prints for later identification. Additionally, high-resolution texture information (micro geometry) can also be sensed by commercial technologies such as profilometers (aka 3D surface profilers) commonly implemented with laser interferometry [39]. These profilometers reconstruct surfaces at a high 3D resolution but cost thousands and are too bulky and heavy to be worn. SpeCam [68] uses a front-facing camera and a screen as a multi-spectral light source on a smartphone to recognize surface materials as important contextual information in mobile computing. SpectroPhone [49] employs the rear camera on a smartphone with warm and cool white flashlight LEDs, yielding photos that can be used to effectively differentiate between a diverse range of materials with SVM. Closer to our project in terms of the wearable form factor is Magic Finger [66], a wearable system that uses a miniature camera worn on the tip of a finger to detect a wide array of surface textures for always-available input (e.g., tap and gestural interactions).

Besides vision-based approaches, electrical sensing has shown promise in recognizing objects. For example, Capacitivo [62] is a capacitive sensing approach that allows fabrics with conductive electrodes to sense objects in contact due to their various dielectric properties. For similar applications, Tessutivo [19] recognizes objects with inductive sensing. Daily objects also exhibit different frequency responses to mechanical excitations, which could be leveraged for recognition, as shown in Knocker [20]. With a wearable form factor, VibEye [40] detects passive hand-held objects by sensing the differences between their resonance profiles. RadarCat [67] detects objects by inferring material properties from the reflected millimeter-wave signals. Higher-frequency RF signals (e.g., Terahertz) have also shown promise in sensing material properties [48]. Other than sensing innate properties of objects, researchers have also tagged objects with RF backscatters to facilitate their detection [5, 13, 31, 43].

While all approaches (including *TextureSight*) recognize hand activities by detecting the presence of side-product signals, e.g., sound, EMI vibration, movements, and RF reflection, *TextureSight* is the first that utilized surface textures as that signal, based on the observation that hand-engaged objects are often of unique surfaces textures. Additionally, we first leveraged laser speckle imaging in a wearable form factor. To optimize the sensing capabilities, we incorporated laser diodes with two wavelengths. We found this unique combination could better unveil nuanced differences between surface textures that allowed us to enhance the accuracy of our system. Moreover, we developed a comprehensive end-to-end signal processing and detection pipeline, ensuring improved efficiency in handling and analyzing the data and enabling a more robust and effective sensing approach for various applications.

### 3 EXPERIENCE SAMPLING STUDY OF HAND-ENGAGED OBJECTS

We did not find datasets surveying hand-engaged objects in everyday settings, complete with labels of objects and their surface materials, which we could use to reveal the unicity and repeatability of the object-texture combination in daily activities. This lack of literature motivated our study in this section. Our hypothesis is that objects routinely engaged with by a user often have unique textures. For example, a backpack will have a different *material* than a mug, a computer mouse, or an orange. Even if objects could be of the same material, they are often made with different *fabrication* techniques (weave patterns and densities of a backpack vs. a couch arm). Even for objects of the same material and fabrication, the chance is high that they could feature different *colors* – another dimension that improves unicity. In this research, we considered all the aforementioned factors (i.e., material, fabrication, and color) as part of texture information that could be leveraged. If the unicity of

the object-texture combination within a user's data is sufficient, texture can serve as equivalent information to objects. It is of note that the unicity does not contradict repeatability – in an extreme hypothetical example of someone using only one object over the span of our study, the unicity will be high for no same texture being found on other objects, while the repeatability will be high also. In practice, participants used multiple objects in this study which lowered the repeatability. However, the unicity remained high if there were no two or more objects sharing the same texture.

### 3.1 Procedure

We conducted a survey study of eleven participants using the experience sampling approach [6]. Our participants were drawn from the student population at a university in North America. We gave participants a smartphone (Samsung Galaxy A01) with an app which we developed to facilitate participants' sampling of objects they engaged with using their hands. Over the course of three days, participants received notifications every 15 minutes from 10:00 a.m. to 9:45 p.m. inquiring about objects with which their dominant hands were engaging. We provided an option to opt-out, which was often used when their hands were not engaged with anything, or when they intentionally wanted to skip the round. After our app guided the participants through the photo taking, it asked them to label the names of the objects. At the end of the study, the participants returned to the researchers for a review process in which researchers reviewed the collected data with the participants to label the material type of surfaces their hands were touching. "Unknown" was labeled when participants had their hands engaged with nothing or skipped the notification. In this study, we recruited 11 participants. We dropped data collected from 4 participants who skipped more than half of the notifications during the study. Being mindful of the difference in activities over weekdays and weekends, we partitioned participants into two groups – those who yielded data only during weekdays (n=4) and those only during weekends (n=3). In total, we collected 716 photos over 11.75 hours from the remaining 7 participants. Among this data, 408 photos (23.5 hours from 4 participants) were collected on weekdays and 308 photos (23.5 hours from 3 participants) on weekends.

### 3.2 Results

Figure 2 shows the collected images from all participants. Note that unknown labels (40.7% of total labels) were marked with crosses. We found 203 objects in total, averaging 29 objects (SD = 8.21) per participant. Across these objects, a total number of 12 material categories were found.

**Unicity** We also realized significant unicity in the study. Upon visual examinations of photos collected in the experience sampling and investigations with participants during the post-sampling walk-through, we realized that it was very unlikely that objects used routinely by a participant would share similar textures (i.e., the same combinations of material, fabrication, and color). Although objects of the same type (e.g., laptops and phones) might have shown identical combinations (e.g., MacBooks of different types could use the same gray, oxidized aluminum), these objects often feature similar functions that do not hinder our overarching goal of object detection, which is to infer user activities. We plotted the relationship between materials and the touched objects in Figure 3 (left). We observed a prominent trend wherein the majority of touched items were composed of plastic, silicone, and metal. Interestingly, these three materials also happen to be the most prevalent substances in our daily lives. For instance, metal is a key component in devices like MacBooks, and plastic is frequently used for items like computer mice.

What adds further intrigue is that these materials are manifested in various forms, such as distinct patterns, colors, and other characteristics. This inherent diversity in attributes contributes to the uniqueness of their combinations and thus the texture of touched objects. In other words, even if objects share the same material composition, the presence of dissimilar colors, patterns, and textures sets them distinctly apart from one another. Through a

comprehensive analysis of each individual participant's touched objects and a thorough review across all seven participants, we did not identify any instances of objects sharing identical textures among all 716 objects in our three-day survey study.

**Repeatability** We also observed significant repeatability of the object-texture combination our participants engaged with over the course of the study. Figure 3 (right) shows this result. We used the object-texture combination as a unique ID for each object our participants showed in their data and cross verified them during the review process. Using these IDs, we were able to calculate how often they were repeated in the collected photos. This also sheds light on the feasibility of the texture-based object/activity detection our work is built upon. High repeatability indicates a higher ratio between performance and calibration effort. In one extreme scenario, if there is only one object that a user wants to detect (e.g., the steering wheel for logging the driving time), the user would only need to calibrate the system with one type of signal. In practice, users often engage a variety of objects but the repeatability test can reveal the feasibility of having carefully weighted detectors for most common signals. This technique is widely used in interactive systems – e.g., having the most frequently used functions at the most obvious locations or with buttons of relatively large sizes. Interestingly, the result from this investigation reveals a long-tail distribution that roughly follows the Pareto Principle <sup>1</sup>. The top 10 most frequently engaged object-texture combinations accounted for around 80% of all samples in the data collection. The percentage is around 60% for the top 3 combinations, revealing high repeatability that texture-based object detection approaches can leverage.

#### 4 SENSING PRINCIPLE

Laser beams are coherent, consisting of light waves that are in phase. When reflected from rough surfaces, these light waves interfere constructively and destructively in 3D, forming bright and dark regions, the union of

<sup>1</sup>Pareto Principle: [https://en.wikipedia.org/wiki/Pareto\\_principle](https://en.wikipedia.org/wiki/Pareto_principle)

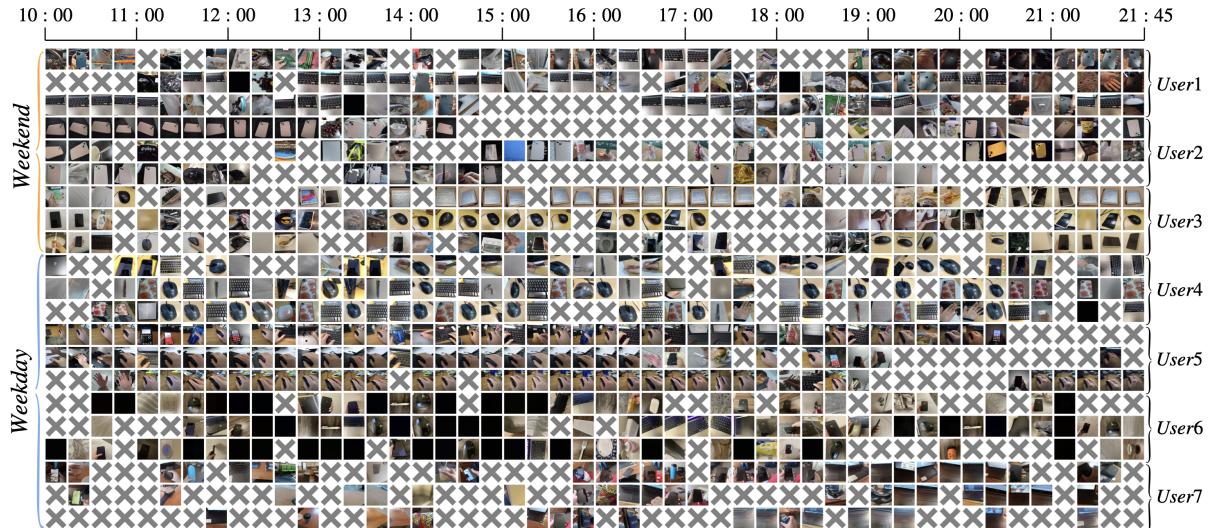


Fig. 2. Photos of hand-engaged objects collected during our experience sampling survey study. "X" image means the participant touched nothing or skipped the notification at that time.

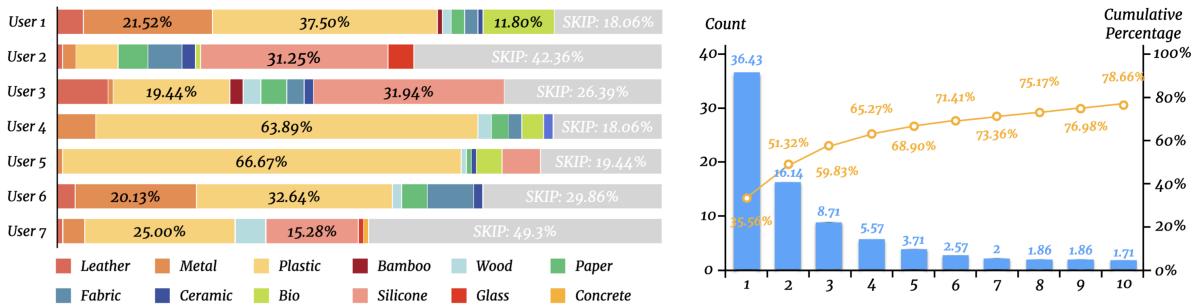


Fig. 3. Left: The distribution of material presence percentages among users. Right: Repeatability test. Top 10 most frequent object-texture combinations – averaged count (blue) and cumulative percentage (orange) in three days.

which is called laser speckles. Prior work has detailed the formation of laser speckles with illustrations [9, 77], which we skipped here for simplicity. The correlation between surface micro-geometries and laser speckles has been discovered in the literature [33, 54, 64]. For example, statistics such as *Contrast*, *Correlation*, *Energy*, and *Homogeneity* calculated using a co-occurrence matrix on speckle images (i.e., an indication of how often pairs of pixels with specific values and in a specified spatial relationship occur in an image) were used as an indicator of surface roughness – a common technique for vision-based systems (see [33] for details). We focused on inherent factors in textures and the imaging systems that affect laser speckle size and shape. In this section, we first introduce the theoretical model we used to correlate the shape of laser speckles with surface micro-geometries, laying the foundation of our sensing principle. We then validate this model with simulations of synthesized speckle images proving the discernibility of speckle patterns generated by different surface textures.

To simulate the laser speckle patterns, we implemented the Gaussian beam as the incident light in MATLAB for simplicity because the Gaussian beam is often the most desirable type of beam provided by a laser source [1]. The rough scatter field is designed by the different surface height function  $h(\alpha, \beta)$ . Equation 1 [60] describes the distribution of the Gaussian beam with wavelength  $\lambda$  on the rough surface after the phase variation induced by microgeometry on the scatter field. In Equation 1,  $k$  is the wavenumber defined by the formula  $\frac{2\pi}{\lambda}$ ;  $\omega_0$  and  $\omega$  are the waist radius and enlightened beam spot radius, respectively. Equation 2 [60] denotes the scattered Gaussian beam propagated for distance  $z$  from the rough surface  $(\alpha, \beta)$  to the observation plane  $(x, y)$  modeled by the Fresnel-Kirchhoff diffraction integral [52]. It is of note that  $(\alpha, \beta)$  corresponds to the coordinates  $(x, y)$  on the surface. We utilize the notation of the  $(x, y)$  coordinates in our observation plane.

$$U_{scatter}(\alpha, \beta) = \frac{\omega_0}{\omega} e^{-(\alpha^2 + \beta^2)(\frac{1}{\omega_0^2} + \frac{ik}{2\rho} - ikz)} \times e^{i\frac{2\pi}{\lambda} h(\alpha, \beta)} \quad (1)$$

$$U_{observe}(x, y) = \frac{e^{ikz}}{i\lambda z} \times \iint_{-\infty}^{\infty} U_{scatter}(\alpha, \beta) e^{\frac{ik}{2z}(\alpha^2 + \beta^2)} e^{-\frac{ik}{z}(x\alpha + y\beta)} d\alpha d\beta \quad (2)$$

#### 4.1 Laser Speckle Patterns

To better explain our observation, we decompose laser speckle patterns into two major aspects – individual and collective, with the individual aspect depicting one laser speckle while the collective aspect depicts interactions between multiple laser speckles. Specifically, the *individual* aspect includes the shape and perimeter of individual

laser speckles. These two factors are affected by the distance between the laser source + imager and the surface  $z$ , the wavelength of the laser  $\lambda$ , and the cross-section of the laser enlightened area on the surface  $\omega$ . We refer to [60] to calculate the mean speckle radius  $R$  (see Equation 3). Generally speaking, the speckle radius increases with the laser wavelength, meaning the distance between the imager and the illuminated area. It is inversely proportional to the size of the illuminated area.

$$R = \frac{\lambda z}{\pi \omega} \quad (3)$$

The *collective* aspect of laser speckles is the spatial pattern of a group of adjacent laser speckles, mostly referring to the spacing between them. These high-level descriptions could not depict the real complexity of laser speckles – since laser speckles have irregular shapes, the "distance" between them cannot be represented with simple mathematical terms nor common wording. We refer to Figure 5 to show laser speckles induced by different real-world surfaces. We will later describe the data collection setup and procedure that generated these images. It is of note that laser speckle patterns consist of both bright regions (the laser speckles) and dark regions due to constructive and destructive interference. Based on this theory, the distance between speckles should equal the speckle size. However, in practice, we found that configurations of the imager also contribute to the distance between speckles. Weak reflections due to darker surfaces or opposing colors of the surfaces and the excitation of the light, could result in more dark regions and thus larger distances between laser speckles (see Figure 5 for examples).

At a micro-scale, a rough surface can be modeled as geometries consisting of micro surfaces at various distances to the imager, manifesting as both the size of speckles and the distance between them. This explains the irregular shapes of speckles. The spatial characteristics of these micro surfaces, defined by surface textures, mostly contribute to the irregularity of speckle shapes, which we hope to infer and use to reveal texture types.

## 4.2 Simulations on Synthesized Surfaces

We simulated surfaces with mathematically simple equations to verify the factors of surfaces that decided laser speckle characteristics, investigating the effects of surface roughness on speckle shape, size, and distribution. We used simple mathematical forms to generate random-roughness surfaces. As in previous simulations, we used Gaussian beams to simulate laser light and Fresnel-Kirchhoff diffraction to model the interference of the backscattered light waves. We first conducted a series of simulations to confirm factors that affect individual and collective aspects of laser speckles using a random rough surface (Gaussian noise,  $\sigma = 50$ ), changing the distance between the surface and the laser-imager bundle (i.e., from 0.5 cm to 1.5 cm at 0.5 cm intervals), the laser wavelength (i.e., 380 nm, 560 nm, 740 nm), and the radius of the illuminated area (i.e., from 50  $\mu\text{m}$  to 150  $\mu\text{m}$  at 50  $\mu\text{m}$  intervals). Figure 4 shows our results. We found the average size of laser speckles increases with distance and wavelength but decreases with enlightened area, which is consistent with our model. All simulations were performed in MATLAB using our custom tool, which performs calculations according to Equations 1, 2, and 3. These findings from synthesized laser speckle images were also confirmed via measurements from real-world surfaces (see examples on Figure 1 and 5) and our empirical observations in the data collection.

## 4.3 Laser Speckles on Real-World Surfaces

To get a sense of laser speckles induced on real-world surfaces, we collected data using an imager with a bare CMOS sensor (2.2  $\mu\text{m}$  pixel size, 2592  $\times$  1944 resolution), from materials including wood, porcelain, plastic, fur, leather, paper, fabric, rubber, foam, metal, and silicone, all with their natural colors or those that are common to find. We disabled all auto-adjusting functions on the camera and configured an exposure time of 20 ms with a gain of 0 dB. The laser-imager bundle was placed 40 mm away from these surfaces with the help of a fixture

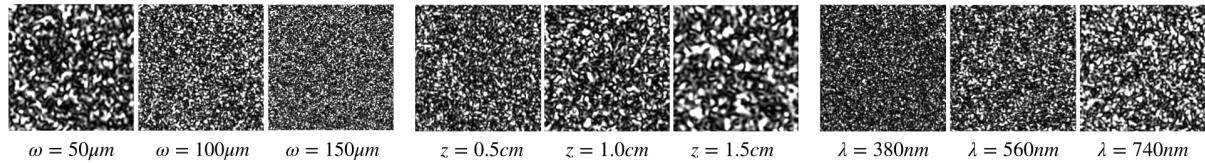


Fig. 4. Synthesized laser speckle images from the simulations with varied parameters of lasers, imagers, and the enlightened surfaces, investigating three key factors – left: illuminated area  $\omega$ , middle: distance between the laser source+imager and the surface  $z$ , right: laser wavelength  $\lambda$ .

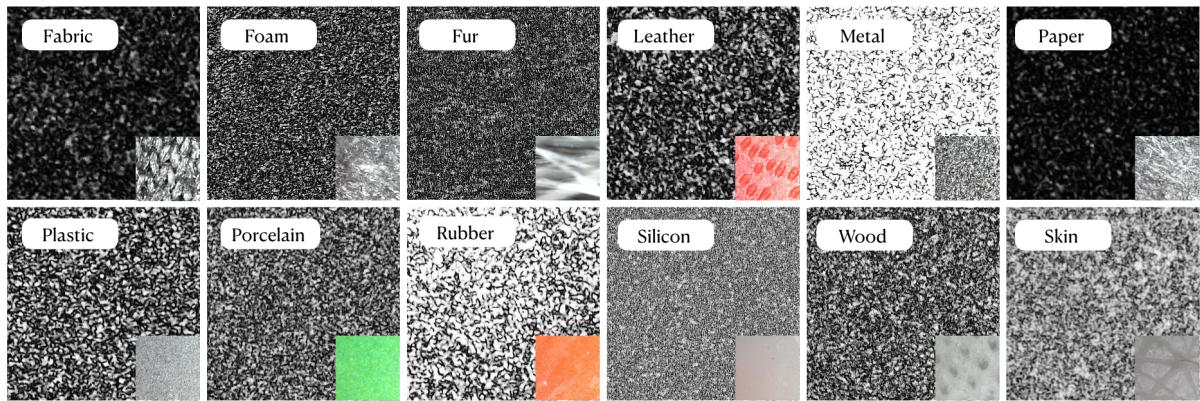


Fig. 5. Gray-scale laser speckle images from 12 real-world surfaces, labeled with material types and microscopic photos.

similar to the enclosure on the final ring prototype. This fixture spaced the laser-imager bundle at a fixed distance away from surfaces when it touched down onto surfaces. Figure 5 shows the collected speckle images labeled with the names of materials and their microscopic images (400x magnification). First, real-world textures exhibit a wider variety of irregular surfaces, each possessing distinct characteristics which are discernible to the naked eye. Additionally, the color of the surface is a major factor in the images of laser speckles. Strong reflections resulted in images with shorter distances between speckles (i.e., stronger overall intensity) due to the innate brightness of the surface color, and its interaction effect with the excitation laser light. This result reinforced our belief in leveraging surface color as part of the texture information in later inference tasks. In the supplementary study, we compared the merits of color with those of micro geometry in regard to the sensing performance.

#### 4.4 Guidelines for Sensor Implementation

This series of simulations and real-world measurements provided guidelines for our sensor implementation. First, an ideal image of laser speckles should show speckles at sufficiently large sizes as their shapes contribute important information to the classification of textures. In this regard, a wearable sensor should *decrease* the illuminated area, *increase* the sensor distance to a surface, and *increase* the laser wavelength (Figure 4). To achieve these desired configurations, we added a mask laser with a tiny hole to the laser diode to reduce the illuminated area, attached the sensor to the far side of the ring prototype against surfaces to achieve the longest possible distance between them both, and chose relatively large laser wavelengths, although the selection of laser wavelength was mainly decided by the experimental results from which we picked the best performers. From real-world measurements, we found that the color of surfaces to be a major factor in the images of laser speckles.

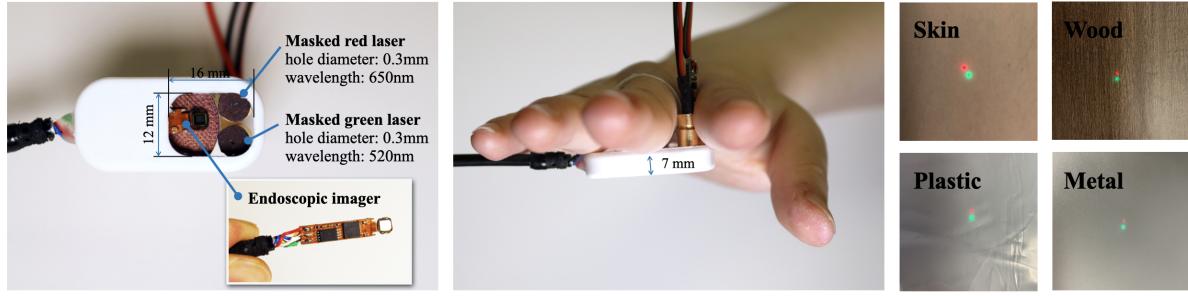


Fig. 6. Ring prototype of *TextureSight*. From left to right: front side (palm side) view of the ring prototype with masked laser diodes and endoscopic imager inside, side view of the ring prototype worn on a user's hand, laser dots landed on object surfaces.

We subsequently decided to use two lasers with different wavelengths to further expose this information for texture detection. Finally, we observed that there were textures, such as paper and fabric which could yield dim speckle images, making it susceptible to noise such as illumination from the ambient environment. To address this, we designed our ring prototype so that the ring body constitutes the enclosure of our sensor, keeping environmental noise from our sensor and the surfaces it senses while maintaining a constant distance between them when the surfaces are in contact with the ring prototype (Figure 6). We document the implementation of this ring prototype in greater detail next.

## 5 TEXTURESIGHT IMPLEMENTATION

### 5.1 Ring Prototype

Our sensor features a ring prototype and the corresponding software. In a nutshell, the ring prototype features two lasers which take turns to illuminate the surface of a hand-engaged object. The reflected light, which manifests as laser speckles, is captured by a bare CCD module. The resulting images are streamed to a PC for signal processing and inference through USB. The PC also runs a Python script that controls the laser toggling and CCD module photo taking.

**5.1.1 Laser.** We used two laser diodes as the excitation laser lights – one green laser of 520 nm wavelength (less than 1mW, Class II) [24] and one red laser of 650 nm wavelength (less than 1mW, Class II) [25], as a simple dual-spectral configuration. It is of note that the output powers mentioned above were all maximum output, the actual consumed power was determined by the actual voltage and current across the lasers. In iterations, we also tried a blue laser diode but eventually removed it for a modest gain in performance and bulky form factor. The laser diodes on the ring prototype were masked with aluminum-coated paper which has a hole of 0.3 mm in the center (Figure 6), limiting the laser beam waist to enlarge the laser speckle size. This design allowed us to capture large speckles that contained rich shape information, preserving the CCD module's pixel density. These two laser diodes were bundled side by side pointing along the direction of the palm (Figure 6). The current implementation has these two lasers fit through the space between the ring and the little finger.

**5.1.2 Repurposed endoscopic imager.** Miniaturized camera systems are commercially available and have been widely used in applications such as endoscopic imaging. We acquired an endoscopic imager circuit board [7] which featured a CCD module with a resolution of 640×480 at 30fps and a signal processor that streamed the formatted pixels to the USB interface. A thin USB cable was tucked along the roots of fingers. Figure 6 shows a

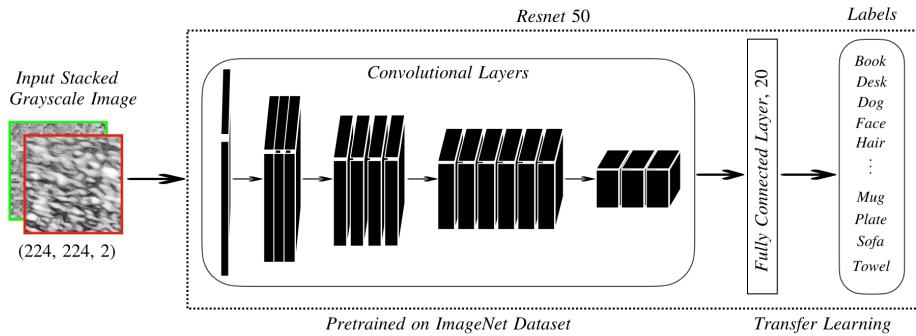


Fig. 7. The ResNet50 model structure with transfer learning as well as the composition of dual speckle images input.

closeup photo of this imager circuit board. This imager was placed inside the ring on the root of the ring finger. The CCD module was angled at around 45° to have the two laser dots around the center of its principal axis for maximum signal strength. We manually selected a region of the highest intensity on the laser speckle image to crop when each laser was on. The size of the cropped images will be introduced later in the software section. This was the only calibration we needed once each ring prototype was made.

**5.1.3 3D-printed ring enclosure.** Figure 6 shows our 3D-printed ring enclosure. The ring could be printed in various sizes to adapt to the wearer's finger width. It was attached to a capsule that measured 36mm (length) by 18mm (width) by 7mm (height). The capsule had a cavity that contained the two lasers and the imager. These components were placed close to the root of the ring finger, ensuring the longest possible distance between the imager and the laser dots, which measured 8.6mm. The bottom of this capsule was hollowed with a 16mm by 12mm opening to allow the emitting and returning light beams to pass through the cavity inside the ring body. This opening was where objects were in contact with the ring prototype. We painted the interior of the capsule black to isolate our sensor from the ambient light which could cause turbulence to the sensing. The enclosed cavity was shielded from the ambient light by our ring enclosure for robustness, and we did not find any effect from illumination in the environment on our sensor's performance.

## 5.2 Software

For toggling the lasers one by one without interference, we used a programmable DC power supply connected to the LAN and then communicated with the PC via socket. The scripts were able to send SCPI commands and retrieve the data back from the DC power supply. We were then able to control the laser in terms of turning different colors on and off at a specific frequency and sequence.

**5.2.1 Signal processing.** The green laser and red laser would be activated and turned off successively, being lit up for a short period of time when responding to the instructions sent from scripts. Note that only one laser would be lit up at one time to avoid mutual interference. While the green or red laser illuminated the surfaces, the CCD module was triggered to capture speckle images synchronously and paired the speckle image from the green laser and the consecutive speckle image from the red laser as an atomic pair, which can be regarded as a signal input for the deep learning model we discuss next. Before a pair of speckle images were fed into the model, the recorded speckle images were cropped by the size of 224×224 pixels, then converted from RGB to grayscale<sup>2</sup>.

<sup>2</sup>Matlab `rgb2gray`:  $I = 0.2989 \times R + 0.5870 \times G + 0.1140 \times B$

**5.2.2 Deep learning inference.** The speckle images captured are abundant in characteristic, minuscule, and unique geometrics features. Convolutional neural networks perform well in learning spatial features with higher effectiveness compared to other techniques [38]. In various deep learning models, residual neural networks (ResNet) [23] are widely used for complex image recognition, offering excellent performance. In our early experiments, ResNet50 performed similarly to other variances and thus was chosen for implementing *TextureSight* for its balanced expressivity and training complexity. We expect other variances of ResNet to work well for our sensing approach, and they might perform better given the resource constraints of application-specific host systems in real-world use cases of *TextureSight*. Our DL model shown in Figure 7 was trained using the *PyTorch* library [42] based on a pre-trained model with the ImageNet dataset [46]. An atomic pair of grayscale speckle images generated by the green laser and red laser illumination were stacked into two layers and fed into the dual-channel model. It is of note that these images were randomly rotated as a common data augmentation technique in the literature [55]. During the training process, a stochastic gradient descent (SGD) optimizer was employed and trained for 60 epochs where the learning rate and batch size equaled 0.0001 and 32, respectively. We chose the state-of-the-art deep learning approach for its superior generalizability and replicability over bespoke ones, lowering the barriers and creating benchmarks for others to leverage our sensor and join forces in wearable laser speckle imaging research.

## 6 EVALUATION

To evaluate *TextureSight*, we conducted a user study including 11 participants. Each participant was instructed to wear the ring prototype and touch/grip various objects and surfaces commonly found in everyday settings, elicited by the experience sampling survey study. The variance in signals induced by different users and their uses is trivial, as the ring prototype is expected to be in contact with objects and surfaces, leaving nearly no leeway. Nonetheless, we incorporated users into our evaluation as they might still induce some unexpected variance due to their varying grip force, hand poses, and locations of touch, all of which – especially those that vary from user to user – are potential factors worthy of investigation.

### 6.1 Setup

11 participants (2 Females, 9 Males, mean age = 25.2) were invited to the study. 20 objects were included in the study (Figure 1). We referred to the objects observed in the survey study and picked those that are common in everyday scenarios, feature a variety of materials, and are available to the source where this research was conducted. Due to time constraints our study to under one hour, only 20 objects were used. However, a thorough investigation with a wider set of objects should be considered in future work for further insights into our proposed sensing technology and its performance. The study was conducted in a lab environment with 19 handheld objects on a table and a large object (i.e., mops) on the ground near the table. Participants were seated in front of the table and could stand if they preferred.

### 6.2 Procedure

There were two data collection sessions for each participant with a 10-minute break in between. In each section, participants were instructed to grip/touch objects as naturally as they would in their everyday uses of these objects. The ring prototype was placed on the ring finger of a participant before the data collection began. After the participant had a grip on or was touching an object, the ring prototype turned on the green laser (the red laser was off) and shone for ~0.3 seconds while 10 photos were collected. The red laser was then turned on (green laser was off) for ~0.3 seconds while another 10 photos were collected. These photos were paired and yielded 10 dual-channel (green+red) laser speckle images (i.e., 10 atomic pairs). The above procedure constituted one *trial*. In total, we collected 10 trials per object in each session, wherein we asked the participant to let go of the object

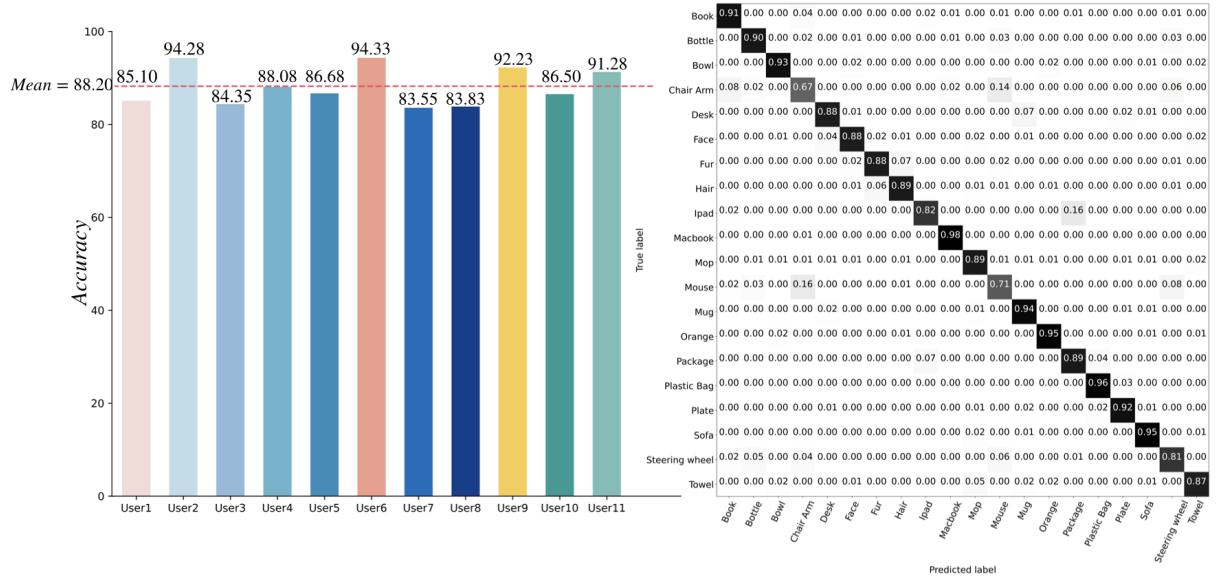


Fig. 8. Accuracy evaluation results for cross-user touching item recognition study – left: classification accuracies among 11 participants in the user study; right: averaged confusion matrix of 20 objects from 11 participants.

and grip/touch it again to start a new trial. We collected 4,000 speckle images (2 sessions  $\times$  10 trials  $\times$  20 objects  $\times$  10 green+red speckle images) for each participant and obtained 44,000 images across 11 users in total.

### 6.3 Results

A cross-user test method was employed to validate object detection performance using dual channel laser speckle. In the user study, we collected data from 11 participants. For each participant, data collected during both sessions of the user study was used as the test dataset; eight participants were randomly selected from the remaining ten participants, excluding for the tested user, and their data was used as the training dataset. The data from the remaining two participants was used as the validation dataset.

Figure 8 shows the classification accuracies among 11 participants and the confusion matrix revealing the recognition accuracies on each object. The average accuracy of the detection of 20 objects detection across 11 participants is 88.20% (SD = 4.13), where precision is 87.33% (SD = 3.57), recall is 85.88% (SD = 4.27), and F1 score is 85.57% (SD = 4.48), which proves the feasibility of wearable laser speckle imaging in detecting hand-engaged objects. Our method also exhibited high detection accuracy (over 90%) for reflective materials like metal (*bottle, MacBook*) and translucent items (*plastic bag*), indicating its adaptability to a wide range of textures and materials. However, challenges were observed in certain object classifications. The confusion matrix indicated lower performance for distinguishing objects *chair arm* and *mouse*, possibly due to their shared plastic material and black color, as well as the similar surface textures. Additionally, objects with curved surfaces like *steering wheel* and *mouse* also negatively affected the classification accuracies. This issue arose because our current ring prototype did not fully conform to touching curved surfaces, leading to variations in laser speckles caused by leaks of ambient light into the ring enclosure and small variations of distance from the sensor to object surfaces.

Returning to the results from our survey study, the long-tail effect should have a significant impact on the actual classification results. Specifically, objects that are more frequently used, if not receiving a sound classification accuracy, could lower the overall classification performance. We believe it also poses an opportunity, for our classifiers to prioritize more frequently used objects (the top 10 objects, which account for 80% cumulative percentage) to achieve a better overall performance and thus user experience. Reducing the number of labels should yield higher classification accuracies (greater than 88.20% on 20 objects) in our work. That being said, we are cautious that recognition of rarely occurring objects should also contribute to user experience, and the tradeoff between overall performance and the recognition of such items is worthy of further investigation in the future.

Due to computational limitations on the smart ring, we opted to collect data and perform offline testing while prioritizing data collection in our user study. However, our data collection pipeline and the test scheme should have yielded insights consistent with real-time tests, with each frame classified and assigned a label independently from neighboring frames and no tested frame included in the training data. To verify this hypothesis, we conducted a series of real-time tests on some example user cases in the following section. The results of these real-time are closely aligned with the outcomes obtained from the user study, affirming the reliability and efficiency of our sensor in real-time object detection.

#### 6.4 Supplemental Studies

The main study focused on the core sensing capability of *TextureSight* – differentiating surfaces of different textures. We also conducted a series of supplemental studies to further our understanding of the proposed sensing technique.

**6.4.1 Dual-laser vs. single-laser configuration.** We investigated the merit of having multiple laser wavelengths (colors) in the inference tasks. Red and green lasers were chosen as they are commonly available and have distinct wavelengths in the visible spectrum. Red lasers typically have a longer wavelength (around 650nm), while green

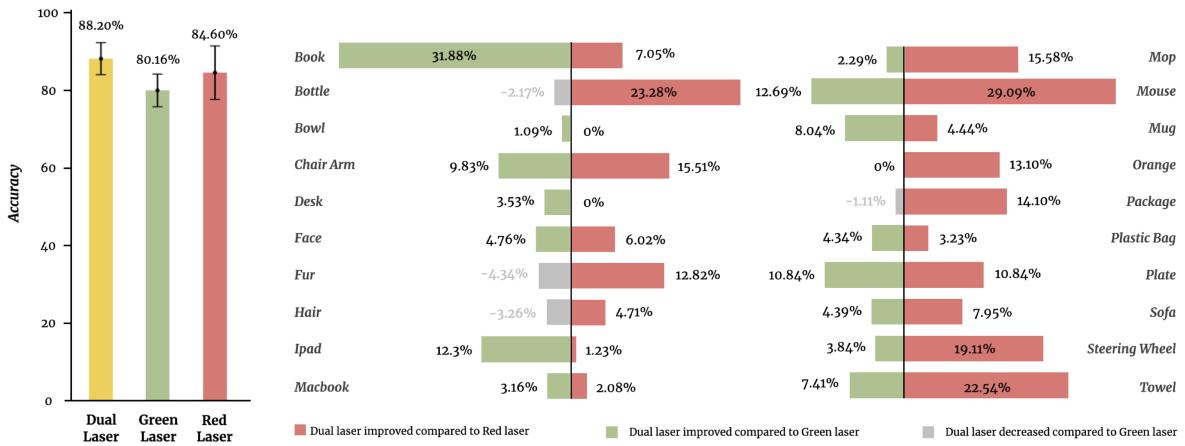


Fig. 9. Evaluation results of touching item recognition with single laser speckle input – Left: Average classification accuracies from a dual and single input, Right: The mean detection accuracy of individual objects was calculated from data collected by 11 users under illumination from either green or red lasers with speckles.

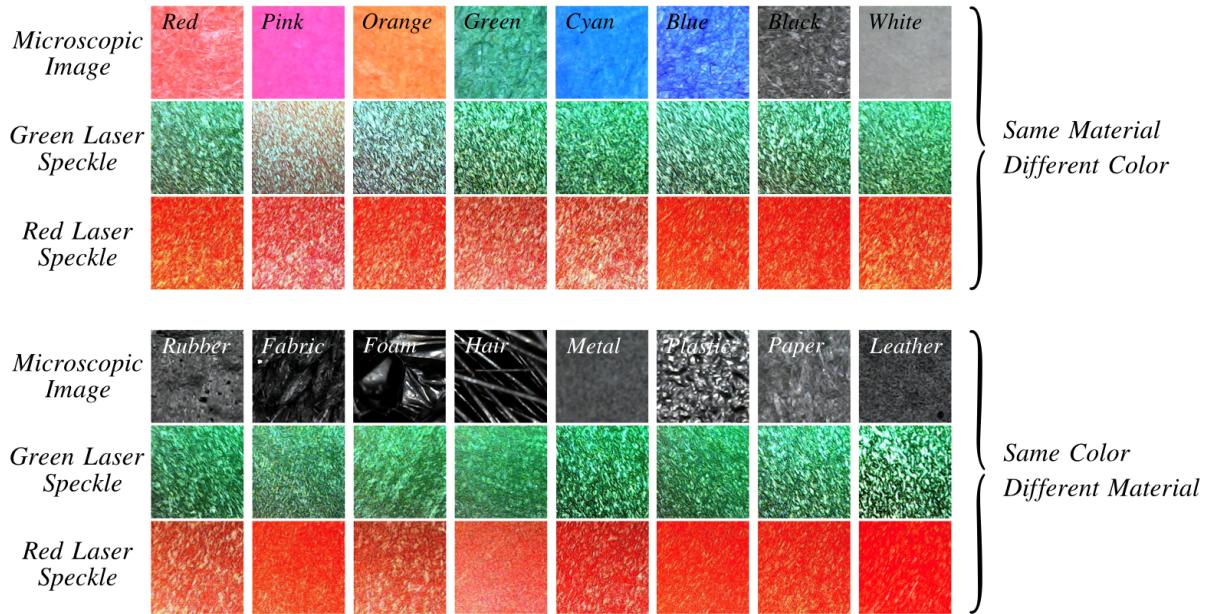


Fig. 10. Top: Microscopic image and speckles illuminated by the green and red laser of papers with 8 different colors. Bottom: Microscopic image and speckles illuminated by the green and red laser of 8 different materials but all in black color.

lasers have a shorter wavelength (around 532nm), resulting in different and complementary laser speckle patterns that maximize the information gain. Moreover, the contrast and visibility of the speckle patterns can be impacted by the laser colors and surface conditions, making it possible for certain materials or surfaces to respond better to one color over the other. In this case, the dual-laser configuration might yield unnecessary noise. To verify this hypothesis, we decided to investigate the performance of touching item classification with single green and red lasers. Specifically, we kept the same backbone of the deep learning structure and the train-test scheme used in our main study, which is the ResNet50 trained for 60 epochs, but varied the input signals – two additional sets of results were reported by using the green laser only and the red laser only. Figure 9 (left) shows the average classification accuracies comparing the performance of dual and single speckle inputs and the confusion matrices resulting from the green-laser-only and red-laser-only configurations. We found that the dual-laser configuration yielded the highest accuracies across 11 participants among the three setups, followed by the red-laser-only configuration, in which the average classification accuracy among 11 participants equaled 84.60% (SD = 6.89). The performance was the worst when only the green laser was used (mean accuracy = 80.16%, SD = 4.21). Furthermore, we investigated the dual laser classification accuracy compared to using red and green lasers on each individual object. Figure 9 (right) shows these breakdown results. When using the red laser, the dual-laser approach consistently outperformed or at least equaled the performance of using the red laser alone, particularly in improving detection accuracy on curved surfaces such as *bottle*, *mouse*, and *steering wheel*. On the other hand, the green laser exhibited greater sensitivity to changes in surface roughness and textures, while a better and more stable detection performance was obtained when employing the dual lasers. This result confirmed our early speculation that different laser wavelengths yielded complementary information about the texture, improving the sensing performance.

**6.4.2 Near-Infrared laser vs. Green/Red laser.** We collected all 20 objects that appeared in the user study from one participant for testing using near-infrared, green, and red lasers. We captured 200 speckle images of each object’s surface, which were acquired under each laser type. These images were then divided into a 7:3 ratio for training and testing purposes. Remarkably, the utilization of solely the near-infrared laser yielded a detection accuracy of approximately 76%, which was notably lower when compared to the performance of the green laser alone (~81%) and the red laser alone (~85%). To highlight disparities in detection accuracy for particular objects, the identification of hair exhibited an accuracy of just 40%, and the detection of a steering wheel fared even worse at a mere 23% with a near-infrared laser. In contrast, the green laser achieved higher accuracy rates for these objects (hair: 92%, steering wheel: 85%), while the red laser also displayed commendable performance (hair: 84%, steering wheel: 62%). Beyond singular laser use, we explored concatenated combinations. The fusion of green and near-infrared lasers resulted in an overall classification accuracy of around 86%, whereas the amalgamation of red and near-infrared lasers achieved an accuracy of roughly 83%. However, all these single laser wavelengths and combinations of wavelengths performed worse than the red and green laser combination (93% accuracy) for that participant.

**6.4.3 Effect of color and micro geometry.** The interaction between excitation laser light and surface color manifests as the intensity of reflection, which affects the distance between laser speckles (as previously discussed in section 4.1). *TextureSight* leverages the color and micro geometry of surfaces both of which are constituent factors that decide surface textures. In our main study, we found that surfaces of similar colors yielded more confusion than others (e.g., mouse and chair arm). To tease out our sensor’s performance along each factor, we conducted two rounds of data collection, each consisting of materials of the same color (i.e., black) and the same texture (i.e., paper). The upper three rows of Figure 10 show the laser speckle images collected from papers of eight different colors, and the bottom three rows show those collected from various surfaces of the same/similar black color. The first row displays the micro geometries of touching surfaces, and speckle images generated under the illumination of the green laser and red laser were placed on the second and third row, respectively. Although the microstructure has slight differences between different-color papers due to the manufacturing and dyeing processes, the speckles look much semblable among each other compared to the different surface textures with a similar black color. Various surfaces, by contrast, offer immense differences in micro geometries which make the speckle images much more distinguishable. Based on this result, we conclude that although both the color and micro geometry of textures contribute to signals that show in laser speckle images, the latter is the dominant factor that decides the shapes of laser speckles, yielding more information for texture (object) recognition.

**6.4.4 Touch segmentation.** Segmenting touch from no touch allows our sensor to preserve power by only making inferences of laser speckle images collected when users’ hands are engaged. In practice, different sensing techniques such as proximity, pressure, and capacitive sensing can be used to segment touch at high accuracies. We also investigated the feasibility of using only laser speckle images in touch segmentation. Figure 11 shows laser speckle images collected when the ring sensor was in contact with various surfaces, and compares them with those collected when the sensor hovered above the surfaces at different distances. Local speckle contrast  $K$  is an important indicator for measuring the blurriness of the speckles, which is defined by the ratio of standard deviation and the mean of the pixel intensities within a predefined area [50]. The value of the local speckle contrast ranged from 0 to 1 (0 meaning the speckles are totally blurry and 1 denoting the speckles are completely static and stable [44]). The laser speckle becomes more salient with the increasing distance between the surfaces and the CCD module.

In our experiments, thresholding the value of the local speckle contrast at 0.37 can segment touch effectively in the presence of ambient light. However, when this experiment was performed without ambient light (Figure 11 last row), we found that the local speckle contrast value alone is not sufficient to segment touch – additional

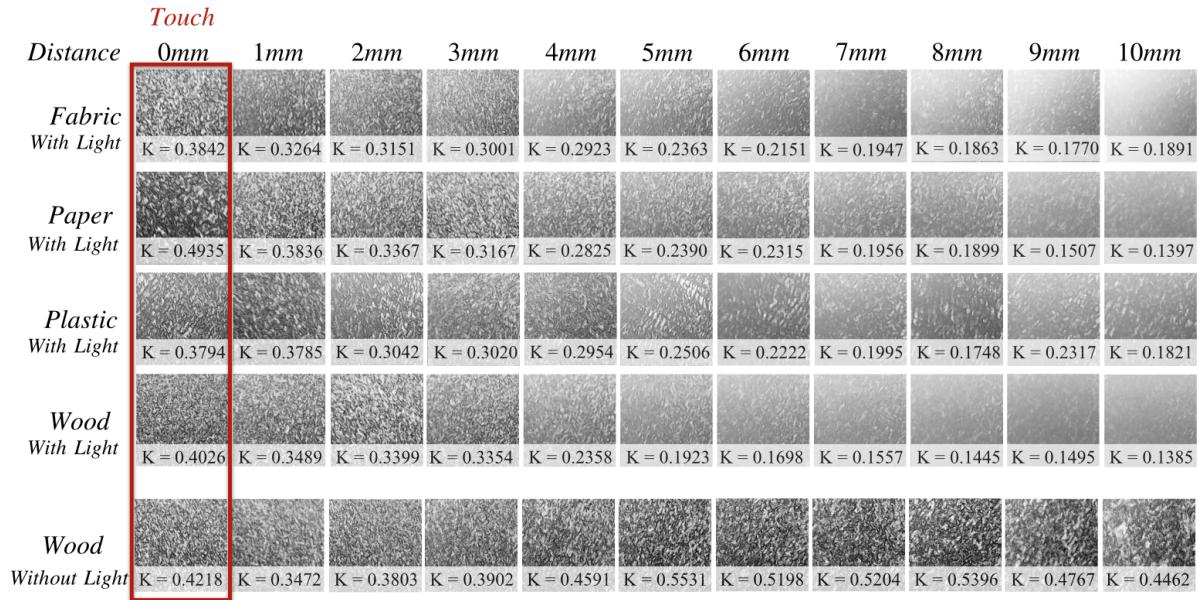


Fig. 11. Laser speckles on four kinds of texture with different distances from 0mm to 10mm with a 1mm interval under two circumstances: the first four rows show the speckles recorded with ambient light; the last row shows the speckles collected under completely dark environments. The local speckle contrast K measures the blurriness of speckles to determine whether the speckle quality for classification, also indicates whether our hand truly touched the object or not.

sensory approaches are needed. Overall, the results indicate that using other sensors (e.g., a force sensor) to turn on the wearable laser speckle imaging sensor should be a more practical and reliable approach in real-world applications. It is of note that although this experiment illustrated that ambient light would influence the quality of speckle imaging, our ring prototype is completely isolated from outside lights except for the laser beam when the user touches the surfaces due to its cavity design.

Finally, we observed frames with blurry speckles, which we term "motional" frames, caused by hand movements when there was loose contact between the smart ring and the object. These motional frames could yield errors in the recognition result, to eliminate such errors, *TextureSight* could use the same aforementioned indicator of blurriness to dismiss all motional frames during touch segmentation and pass only high-quality laser speckle images when there is solid contact between the smart ring and the engaged objects. Additional assistive sensors able to measure contact force and lateral movements could also help identify motional frames. The current implementation of our system does not include this elimination of motional frames, which we acknowledge as a limitation and plan to implement in future iterations of the systems.

## 7 EXAMPLE USE CASES

To demonstrate *TextureSight*, we developed a series of exemplar use cases spanning across activity logging, food journal, smart environment automation, workflow assistance, authentication, ID tags, and education (see also the Video Figure). We also show a clustering technique to enable "few-shot" use experiences that lower the burden of user labeling. Note that all of the following scenes are powered by the aforementioned detection pipeline in the paper.



Fig. 12. *TextureSight* enables logging of activities in a kitchen example: A) using the espresso machine (touching its portafilter handle), B) consuming paper towels, and C) taking out food from the fridge.

**Activity Logging** Most activities have their regular and specific indicator objects. For instance, people usually touch keyboards and mice when they are studying or working, whereas they touch knives, forks, and spoons while eating. In this case, detected textures could be translated into user activities, the logging of which, could power applications of personal informatics. In this example, *TextureSight* could keep track of which type of activities its wearer has been doing and for how long (Figure 12). Similarly, *TextureSight* could recognize surfaces of fruits and vegetables to facilitate food journals (Video Figure).

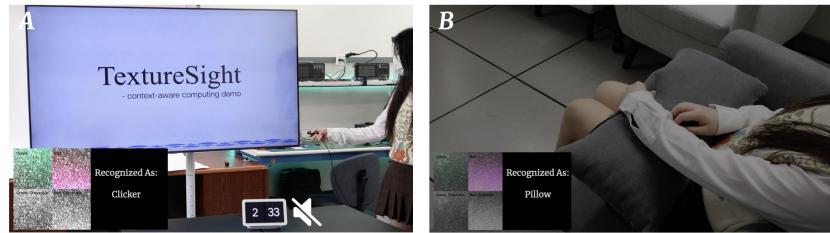


Fig. 13. *TextureSight* could power context-aware applications such as home automation: A) speaker is muted during presentations, and B) light is dimmed when a pillow is grasped.

**Smart Environment Automation** With *TextureSight*, smart environments could better assist users with tasks after acquiring their activity information, enabling applications of context-aware computing (Figure 13). For example, speakers in the room could be muted once *TextureSight* detects a user's grasp of a presentation remote as a strong indication of someone giving a presentation. Similarly, in a home environment, lights could be automatically dimmed once *TextureSight* detects a user's grasp of a pillow. The smart ring can automatically adjust living atmospheres such as light and temperature based on the behavior and state of the users as perceived by the ring.

**In-Situ Work Assistance** Hand tools are often made of different materials or the same materials with different colors and surface textures (Figure 14). *TextureSight* could detect which tool is currently being used by the user to allow smart devices, such as a smart speaker, to provide in-situ instructions for training purposes, or in DIY projects, like the furniture assembly example shown here. Specifically, *TextureSight* loads the model for tool recognition once the user touches the surface of the furniture piece. It could recommend to the user the right tool for the task (e.g., drilling a screw into wood with a hand drill instead of a screwdriver).

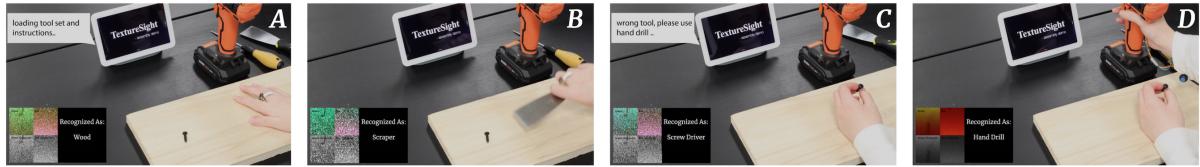


Fig. 14. *TextureSight* provides pinpoint instructions in-situ, loading tool set and instructions once the surface of furniture is detected (A), alerting the user of misuse of tool (B,C,D).



Fig. 15. *TextureSight* could load application-specific models (texture sets) for improved performance and eliminate the need for collecting personal data. In this field trip example, publicly available models of common plants and their parts could be readily loaded and used for educational purposes.

**Plant Recognition for Education** Textures on plant leaves are unique features that could play an important role in plant classification. With the support of the huge plant database, *TextureSight* could record and recognize leaves, then educate its users about the plants. We created a simplified demonstration of this use case (Figure 15). A student on a field trip could use *TextureSight* as an always-available sensor and intelligent tutor to recognize plants for educational purposes.

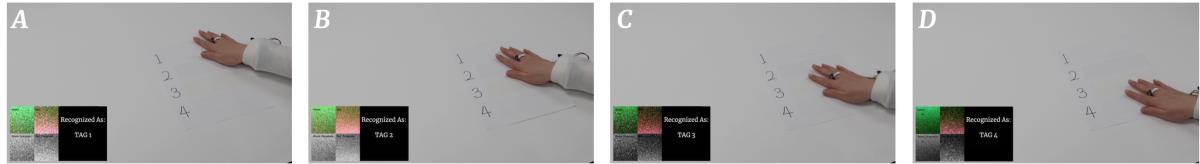


Fig. 16. *TextureSight* enables texture-based information to be encoded into tags that are visually similar for identification purposes.

**Texture-based ID Tags** Beyond detecting textures as passive features on object surfaces, we could also explicitly assign them ID tags as a way of encoding information. As a demonstration of feasibility, we used white tapes commonly found on the market as proof-of-concept texture ID tags (Figure 16). These pieces of tape look visually similar but are of different textures, allowing a user to assign IDs to different objects. These texture-based ID tags could be used to facilitate the recognition of textureless surfaces or to disambiguate objects of identical surface textures. Additionally, they allow users to explicitly tag objects of interest to facilitate smart environment applications.

**Authentication** The smart ring in which *TextureSight* resides could function as a key in authentication applications (Figure 17). In this smart lock example, *TextureSight* could authenticate its user (i.e., unlock the door) once



Fig. 17. *TextureSight* turns a smart ring into keys in authentication applications by detecting its contact with locks with unique textures.

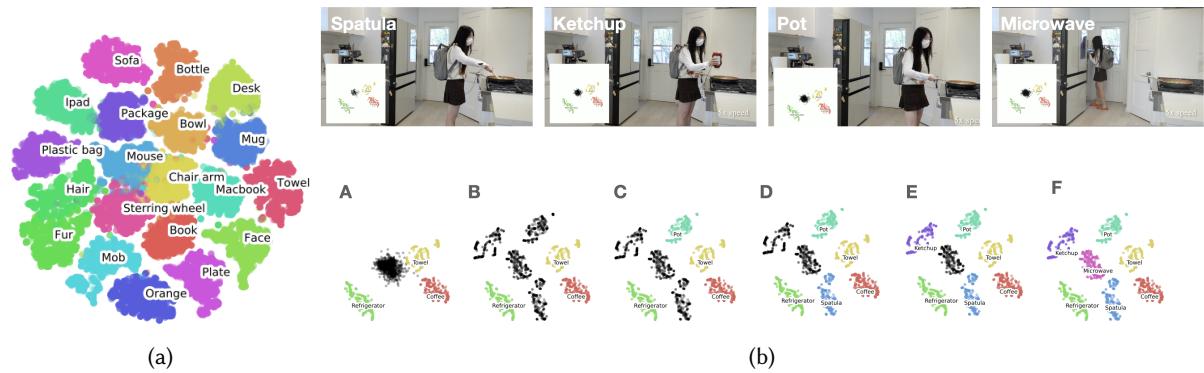


Fig. 18. (a) Speckle clusters of 20 objects in the user study. (b) Top: the speckle images collected when the user touched different objects without labeling; Bottom: The illustration for newly collected speckle data clustering process (A → B) and the one-shot labeling process (C → F).

it detects its user's grasp on the handle of the smart lock, with the smart lock knowing that *TextureSight* is worn on its user's finger. Such an interaction scenario is akin to authentication enabled by NFC but it exists without the need for the underlying wireless communication transceivers which could be expensive and incompatible with metal surfaces.

**"Few-Shot" Experience Through Unsupervised Clustering** *TextureSight* could already recognize a wide array of textures "out of box" without the need for collecting personal data. Even for collecting personal data, our work is compatible with the unsupervised clustering technique inspired by prior work (e.g., Wu et al. [61]) to lower the effort of user labeling. As discussed previously, we employed a pre-trained ResNet50 model to extract features on speckle images collected from 20 different surfaces that appeared in the user study, then mapped the extracted features from high dimension to a plane through a nonlinear dimensionality reduction method. Figure 18 shows the cluster results among 20 objects and demonstrates the separability of speckle patterns acquired from diverse textures. This observation indicates that it is possible to cluster the unlabelled speckle images based on the features extracted by a pre-trained model, then label each cluster simultaneously at a later date, lowering user burden in labeling.

Returning to our previous kitchen example (Figure 18), *TextureSight* could cluster data points from unseen surfaces first using deep learning featurization and t-distributed stochastic neighbor embedding from seen surfaces, then

wait for the user to label these data points later, after clusters are formed. This labeling only needs to be done once for each cluster.

## 8 DISCUSSIONS

### 8.1 Establishment of Dataset

One limitation of our technique comes with the assumption that an object correlates with one or several specific textures. This assumption would require the establishment of a personal dataset. We acknowledge that building a personalized dataset is an interaction overhead that might undermine the practicality of any sensing technique. However, many textures are universally unique to their objects and thus could be trained in the factory to enable “out of box” uses. Such textures can be found on surfaces of common everyday objects such as fruits, vegetables, and plants. Additionally, we could explicitly assign textures to objects in a labeling process by attaching stickers with textures our proposed technique could have been trained in the factory.

For establishing a personalized dataset, prior work has shown a promising technique to minimize user labor in labeling. A self-supervised clustering algorithm was designed [61] to cluster unlabelled speckle images cumulatively before asking users for labeling (once, and for all speckle images). As we show in the example use cases (Section 7), we clustered data points from unseen surfaces first using deep learning featurization and t-distributed stochastic neighbor embedding from seen surfaces, then waited for the user to label these data points later, after clusters are formed. Figure 18 shows the clusters of unlabeled data points. This labeling only needs to be done once (see also the Video Figure).

### 8.2 Power Consumption

Low power consumption is one of the keys to the success of wearable sensing. We investigated power consumption spent on the sensing without the system power consumption overhead. Our sensor consumes energy mostly for two main parts: 1) the two laser diodes, and 2) the CCD sensor. The neglected power consumption includes chip-to-chip communications (e.g., from co-processor to MCU) and the power regulations (for laser diodes and the CCD module). Despite the already-low power consumption of our laser diodes, it can be even lower with custom laser diodes that output only narrow laser beams. The output power of our green laser is 1 mW with an output aperture equal to 4.3 mm and the red laser is 1 mW with an output aperture equal to 3 mm. Considering the conversion efficiencies of these two lasers after masking the beams with 0.3mm holes, according to their datasheets and assuming the linear relationship between the power consumption and enlightened area, the estimated aforementioned output powers translates to 0.005 mW and 0.01 mW energy consumption for the hypothetical custom laser diodes, respectively.

The 640x480 CCD module consumes ~100mW and we suspect it can be lower with potentially lower duty cycles enabled by low-power touch segmentation techniques (e.g., laser proximity sensors) that turn on the CCD module only at touch events. Additionally, we only used 224x224 pixels on the 640x480 CCD. In theory, a custom-made CCD of *TextureSight* could consume only 16.33mW or even less, assuming a linear relationship between the number of pixels and the power consumption of CCD and with its power supply piggybacking on existing power rails of wearables. Note that we omit communication (e.g., inter-chip comm.) and computing for their high-energy efficiencies within commercial products. The hypothetical low-end power consumption of main components (one CCD and two lasers) in our sensing approach is 16.35 mW, corresponding to 5.7 hours of continuous operation on a 22 mAh Li-Po battery (assuming a 4.2V supply voltage), which can be found on a commercial smart ring product <sup>3</sup>. Although these stats show feasibility, they could be further improved with smart scheduling (e.g.,

<sup>3</sup>Our Ring: <https://support.ouraring.com/hc/en-us/articles/360025428394-Product-Safety-Use>

sense once only at the touch events – transitioning from no touch to touch) and future optimization of hardware (e.g., use ultra-low-power CCD, or a CMOS sensor).

### 8.3 Miniaturization

Our current implementation is bulky without much effort in miniaturization. The heat sinks of the lasers we used take up ~65.3% of the size of our ring prototype and can be eliminated for the low-power and low-duty-cycle nature of our proposed sensing approach in commercialization. Better ergonomics should and can be achieved with custom surface-mount laser diodes and an integrated PCB with all the components built in. Though we suspect that the sensing performance will not become significantly higher because of the improved ergonomics, improvements along this line could improve comfort, especially for long-term use. Furthermore, our sensor requires a cavity inside the capsule for the incident and reflected laser beams. Though small, this cavity might take up valuable space on wearables that could have been preserved for batteries.

### 8.4 Framerate

We have also experimented with a simultaneous laser illumination approach, hoping to achieve a higher framerate by extracting laser speckle images induced by the green and red lasers from the corresponding channels of the RGB image. We used the 20 objects in the main study and followed the same procedure but with two different laser illumination approaches. We invited one participant to this experiment. The average accuracy was 93.48% when the green and red lasers were switched on and off alternately. This accuracy dropped to 88.92% with the simultaneous laser illumination approach due to the imperfection of the green and red filters on the CCD sensor – the green laser could induce a charge on red pixels and vice versa, resulting in interference and thus the degradation of information revealed by the two lasers. This result shows promise nonetheless, and thus we expect CCD modules of higher build quality and filter performance to resolve this interference issue and yield a higher framerate.

### 8.5 Alternative Form Factor

Our main study showed a relatively modest performance on objects with curved surfaces due to the limitation that comes with the rigidity of the ring prototype, especially the capsule part for introducing variances in distance and thus producing signals that are more challenging to infer. We hope to address this limitation with flexible form factors that could comply with the curvature of surfaces. Additionally, we suspect that wearable laser speckle imaging could be useful in potentially acquiring contextual information on a wide array of instrumentation locations including headsets, garments, smartwatches, and shoes, due to its advantages of being accurate, free of a camera lens, and potentially low in power.

### 8.6 Privacy Concerns for Smart Ring

The wearable smart ring equipped with only a CCD sensor without lens presents an intriguing solution to address privacy concerns related to traditional vision-based wearable devices. By eliminating lenses, our smart rings could only image laser speckles, which are induced by laser beams that are always in focus while not being capable of imaging the background from the ambient illumination, ensuring that no real images or visual data that might contain privacy-sensitive information are collected even if our smart ring is accidentally activated in the air. Furthermore, no images are captured when the lasers are deactivated or when potential external sensors detect no contact, which further reduces the duty cycle of imaging and protects user privacy.

## 8.7 Safety Concerns of Lasers

We selected low-power laser diodes with an output power of less than 1mW, which poses minimal risk to the human eyes and skin and adheres to Class II laser safety classification [18]. In summary, while direct eye exposure to Class II lasers should be avoided, they are relatively safe because the human blink reflex (aversion response) to bright light typically limits the exposure to a safe duration. Additionally, the mask lasers we put in front of these diodes further enhance safety – the less-than-1mW beams are masked with a 0.3mm layer of light-proof paper, preventing excess direction exposure to the laser light outside the device. Finally, activation of the laser diodes is synchronized with the imager, which further reduces the duty cycle and thus the output power of the lasers to improve the safety of use.

## 9 CONCLUSION

We proposed *TextureSight*, a wearable laser speckle imaging sensor that recognizes objects based on their textures. Our sensor worked best for objects which are routine to a user and among which correlations between textures and objects are strong. We conducted an experience-sampling study to investigate the unicity and repeatability of object-texture combinations. Theoretical models were created and validated with simulations to ground our sensing principle. We implemented a ring prototype with two lasers and a bare CCD imager. State-of-the-art Deep Neural Net techniques were used to power the inference pipeline. We demonstrated an accuracy of 88.2% across 20 objects and conducted a series of supplementary studies to tease out our sensor’s performance envelope. We believe our system uniquely demonstrates the potential of wearable laser speckle imaging in detecting routine objects, and it could lead to practical sensing solutions which are much needed in today’s wearable devices.

## REFERENCES

- [1] Javier Alda. 2003. Laser and Gaussian beam propagation and transformation. *Encyclopedia of optical engineering* 999 (2003).
- [2] Rawan Alharbi, Sougata Sen, Ada Ng, Nabil Alshurafa, and Josiah Hester. 2022. ActiSight: Wearer Foreground Extraction Using a Practical RGB-Thermal Wearable. In *2022 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 237–246.
- [3] Daniel Ashbrook, Patrick Baudisch, and Sean White. 2011. Nenya: subtle and eyes-free mobile input with a magnetically-tracked finger ring. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2043–2046.
- [4] David A Boas and Andrew K Dunn. 2010. Laser speckle contrast imaging in biomedical optics. *Journal of biomedical optics* 15, 1 (2010), 011109.
- [5] Michael Buettner, Richa Prasad, Matthai Philipose, and David Wetherall. 2009. Recognizing daily activities with RFID-based sensors. In *Proceedings of the 11th international conference on Ubiquitous computing*. 51–60.
- [6] Mihaly Csikszentmihalyi and Reed Larson. 2014. Validity and reliability of the experience-sampling method. In *Flow and the foundations of positive psychology*. Springer, 35–54.
- [7] Dengofng. 2022. USB Type-C Endoscope, Product Information. [https://www.amazon.com/Dengofng-Endoscope-Waterproof-Inspection-Compatible/dp/B09MRYHVF/ref=sr\\_1\\_14?keywords=Dengofng&qid=1670359259&sr=8-14](https://www.amazon.com/Dengofng-Endoscope-Waterproof-Inspection-Compatible/dp/B09MRYHVF/ref=sr_1_14?keywords=Dengofng&qid=1670359259&sr=8-14) Last accessed 7 Dec 2022.
- [8] B Dhanasekar, N Krishna Mohan, Basanta Bhaduri, and B Ramamoorthy. 2008. Evaluation of surface roughness based on monochromatic speckle correlation using image processing. *Precision Engineering* 32, 3 (2008), 196–206.
- [9] Mustafa Doga Dogan, Steven Vidal Acevedo Colon, Varnika Sinha, Kaan Aksit, and Stefanie Mueller. 2021. SensiCut: Material-Aware Laser Cutting Using Speckle Sensing and Deep Learning. In *The 34th Annual ACM Symposium on User Interface Software and Technology*. 24–38.
- [10] Mustafa Doga Dogan, Faraz Faruqi, Andrew Day Churchill, Kenneth Friedman, Leon Cheng, Sriram Subramanian, and Stefanie Mueller. 2020. G-ID: identifying 3D prints using slicing parameters. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [11] Andrew K Dunn. 2012. Laser speckle contrast imaging of cerebral blood flow. *Annals of biomedical engineering* 40, 2 (2012), 367–377.
- [12] Sune Bro Duun, Rasmus G Haahr, Karen Birkelund, and Erik V Thomsen. 2009. A ring-shaped photodiode designed for use in a reflectance pulse oximetry sensor in wireless health monitoring applications. *IEEE Sensors Journal* 10, 2 (2009), 261–268.
- [13] Kenneth P Fishkin, Matthai Philipose, and Adam Rea. 2005. Hands-on RFID: Wireless wearables for detecting use of objects. In *Ninth IEEE International Symposium on Wearable Computers (ISWC’05)*. IEEE, 38–41.

- [14] Kevin R Forrester, John Tulip, Catherine Leonard, Cody Stewart, and Robert C Bray. 2004. A laser speckle imaging technique for measuring tissue perfusion. *IEEE transactions on biomedical engineering* 51, 11 (2004), 2074–2084.
- [15] Masaaki Fukumoto and Yoshinobu Tonomura. 1997. “Body coupled FingerRing” wireless wearable keyboard. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. 147–154.
- [16] Bogdan-Florin Gheran, Jean Vanderdonckt, and Radu-Daniel Vatavu. 2018. Gestures for smart rings: Empirical results, insights, and design implications. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 623–635.
- [17] Bogdan-Florin Gheran and Radu-Daniel Vatavu. 2020. From controls on the steering wheel to controls on the finger: using smart rings for in-vehicle interactions. In *Companion Publication of the 2020 ACM Designing Interactive Systems Conference*. 299–304.
- [18] Anders Glansholm. 2002. New laser classes. *RADIATION PROTECTION IN THE 2000S—THEORY AND PRACTICE* (2002), 232.
- [19] Jun Gong, Yu Wu, Lei Yan, Teddy Seyed, and Xing-Dong Yang. 2019. Tessutivo: Contextual interactions on interactive fabrics with inductive sensing. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 29–41.
- [20] Taesik Gong, Hyunsung Cho, Bowon Lee, and Sung-Ju Lee. 2019. Knocker: Vibroacoustic-based object recognition with smartphones. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 3, 3 (2019), 1–21.
- [21] Michael A Grandner, Zohar Bromberg, Zoe Morrell, Arnulf Graf, Stephen Hutchison, and Dustin Freckleton. 2022. Performance of a Multisensor Smart Ring to Evaluate Sleep: In-Lab and Home-Based Evaluation Relative to Polysomnography and Actigraphy: Importance of Generalized Versus Personalized Scoring. *medRxiv* (2022), 2021–12.
- [22] Jeremy Gummesson, Bodhi Priyantha, and Jie Liu. 2014. An energy harvesting wearable ring platform for gestureinput on surfaces. In *Proceedings of the 12th annual international conference on Mobile systems, applications, and services*. 162–175.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [24] Quarton inc. 2022. Direct Green Laser Diode Module. <https://www.quarton.com/uploadfiles/1028/product/Green-dot-laser-modules/Economic/VLM-520-74-series/Green-Dot-Laser-Module-VLM-520-74-manual.pdf> Last accessed 7 Dec 2022.
- [25] Quarton inc. 2022. Economical Laser. <https://www.quarton.com/uploadfiles/1028/product/Red-dot-laser-modules/Economic/VLM-650-04-series/Red-Dot-Laser-Module-VLM-650-04-manual.pdf> Last accessed 7 Dec 2022.
- [26] Kensei Jo, Mohit Gupta, and Shree K Nayar. 2015. Spedo: 6 dof ego-motion sensor using speckle defocus imaging. In *Proceedings of the IEEE International Conference on Computer Vision*. 4319–4327.
- [27] Wolf Kienzle, Eric Whitmire, Chris Rittaler, and Hrvoje Benko. 2021. Electroring: Subtle pinch and touch detection with a ring. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [28] Gierad Laput and Chris Harrison. 2019. Sensing fine-grained hand activity with smartwatches. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [29] Gierad Laput, Robert Xiao, and Chris Harrison. 2016. Viband: High-fidelity bio-acoustic sensing using commodity smartwatch accelerometers. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 321–333.
- [30] Gierad Laput, Chouchang Yang, Robert Xiao, Alanson Sample, and Chris Harrison. 2015. Em-sense: Touch recognition of uninstrumented, electrical and electromechanical objects. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. 157–166.
- [31] Hanchuan Li, Eric Brockmeyer, Elizabeth J Carter, Josh Fromm, Scott E Hudson, Shwetak N Patel, and Alanson Sample. 2016. Paperid: A technique for drawing functional battery-free wireless interfaces on paper. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 5885–5896.
- [32] Hyunchul Lim, Jungmin Chung, Changhoon Oh, SoHyun Park, and Bongwon Suh. 2016. OctaRing: examining pressure-sensitive multi-touch input on a finger ring device. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 223–224.
- [33] Rong-Sheng Lu, Gui-Yun Tian, Duke Gledhill, and Steve Ward. 2006. Grinding surface roughness measurement based on the co-occurrence matrix of speckle pattern texture. *Applied Optics* 45, 35 (2006), 8839–8847.
- [34] Takuya Maekawa, Yasue Kishino, Yutaka Yanagisawa, and Yasushi Sakurai. 2012. WristSense: wrist-worn sensor device with camera for daily activity recognition. In *2012 IEEE International Conference on Pervasive Computing and Communications Workshops*. IEEE, 510–512.
- [35] Takuya Maekawa, Yutaka Yanagisawa, Yasue Kishino, Katsuhiko Ishiguro, Koji Kamei, Yasushi Sakurai, and Takeshi Okadome. 2010. Object-based activity recognition with heterogeneous sensors on wrist. In *International Conference on Pervasive Computing*. Springer, 246–264.
- [36] Guillaume Mahé, Anne Humeau-Heurtier, Sylvain Durand, Georges Leftheriotis, and Pierre Abraham. 2012. Assessment of skin microvascular function and dysfunction with laser speckle contrast imaging. *Circulation: Cardiovascular Imaging* 5, 1 (2012), 155–163.
- [37] Md Shaad Mahmud, Honggang Wang, and Hua Fang. 2018. SensoRing: An integrated wearable system for continuous measurement of physiological biomarkers. In *2018 IEEE International Conference on Communications (ICC)*. IEEE, 1–7.
- [38] Adugna G Mullissa, Claudio Persello, and Valentyn Tolpekin. 2018. Fully convolutional networks for multi-temporal SAR image classification. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 6635–6638.
- [39] Nanovea. 2022. Profilometer. <https://nanovea.com/profilometers/> Last accessed 7 Aug 2022.

[40] Seungjae Oh, Gyeore Yun, Chaeyong Park, Jinsoo Kim, and Seungmoon Choi. 2019. VibEye: Vibration-mediated object recognition for tangible interactive applications. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–12.

[41] Farshid Salemi Parizi, Eric Whitmire, and Shwetak Patel. 2019. Auraring: Precise electromagnetic finger tracking. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 4 (2019), 1–28.

[42] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).

[43] Matthai Philipose, Kenneth P Fishkin, Mike Perkowitz, Donald J Patterson, Dieter Fox, Henry Kautz, and Dirk Hahnel. 2004. Inferring activities from interactions with objects. *IEEE pervasive computing* 3, 4 (2004), 50–57.

[44] Dmitry D Postnov, Xiaojun Cheng, Sefik Evren Erdener, and David A Boas. 2019. Choosing a laser for laser speckle contrast imaging. *Scientific reports* 9, 1 (2019), 1–6.

[45] M Roustit, C Millet, S Blaise, B Dufournet, and JL Cracowski. 2010. Excellent reproducibility of laser speckle contrast imaging to assess skin microvascular reactivity. *Microvascular research* 80, 3 (2010), 505–511.

[46] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115, 3 (2015), 211–252.

[47] Munehiko Sato, Shigeo Yoshida, Alex Olwal, Boxin Shi, Atsushi Hiyama, Tomohiro Tanikawa, Michitaka Hirose, and Ramesh Raskar. 2015. Spectrans: Versatile material classification for interaction with textureless, specular and transparent surfaces. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 2191–2200.

[48] Scott Schecklman, Lisa M Zurk, Samuel Henry, and Gabriel P Kniffin. 2011. Terahertz material detection from diffuse surface scattering. *Journal of Applied Physics* 109, 9 (2011), 094902.

[49] Maximilian Schrapel, Philipp Etgeton, and Michael Rohs. 2021. SpectroPhone: Enabling Material Surface Sensing with Rear Camera and Flashlight LEDs. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–5.

[50] Janaka Senarathna, Abhishek Rege, Nan Li, and Nitish V Thakor. 2013. Laser speckle contrast imaging: theory, instrumentation and applications. *IEEE reviews in biomedical engineering* 6 (2013), 99–110.

[51] Soroush Shahi, Rawan Alharbi, Yang Gao, Sougata Sen, Aggelos K Katsaggelos, Josiah Hester, and Nabil Alshurafa. 2022. Impacts of Image Obfuscation on Fine-grained Activity Recognition in Egocentric Video. In *2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*. IEEE, 341–346.

[52] CJR Sheppard and M Hrynevych. 1992. Diffraction by a circular aperture: a generalization of Fresnel diffraction theory. *JOSA A* 9, 2 (1992), 274–281.

[53] Yi Chang Shih, Abe Davis, Samuel W Hasinoff, Frédéric Durand, and William T Freeman. 2012. Laser speckle photography for surface tampering detection. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 33–40.

[54] Motochika Shimizu, Hiroshi Sawano, Hayato Yoshioka, and Hidenori Shinno. 2015. Multi-dimensional assessment of precision machined surface texture based on laser speckle pattern analysis. *Procedia CIRP* 33 (2015), 251–256.

[55] Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of big data* 6, 1 (2019), 1–48.

[56] Brandon M Smith, Pratham Desai, Vishal Agarwal, and Mohit Gupta. 2017. CoLux: Multi-object 3d micro-motion analysis using speckle imaging. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.

[57] Wei Sun, Franklin Mingzhe Li, Congshu Huang, Zhenyu Lei, Benjamin Steeper, Songyun Tao, Feng Tian, and Cheng Zhang. 2021. ThumbTrak: Recognizing Micro-finger Poses Using a Ring with Proximity Sensing. In *Proceedings of the 23rd International Conference on Mobile Human-Computer Interaction*. 1–9.

[58] Hsin-Ruey Tsai, Min-Chieh Hsiu, Jui-Chun Hsiao, Lee-Ting Huang, Mike Chen, and Yi-Ping Hung. 2016. TouchRing: subtle and always-available input using a multi-touch ring. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. 891–898.

[59] Edward J Wang, Tien-Jui Lee, Alex Mariakakis, Mayank Goel, Sidhant Gupta, and Shwetak N Patel. 2015. Magnisense: Inferring device interaction using wrist-worn passive magneto-inductive sensors. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 15–26.

[60] Will J Warren, Erik A Moro, Matthew E Briggs, and Eric B Flynn. 2014. Simulating translation-induced laser speckle dynamics in photon Doppler velocimetry. *Applied Optics* 53, 21 (2014), 4661–4668.

[61] Jason Wu, Chris Harrison, Jeffrey P Bigham, and Gierad Laput. 2020. Automated Class Discovery and One-Shot Interactions for Acoustic Activity Recognition. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.

[62] Te-Yen Wu, Lu Tan, Yuji Zhang, Teddy Seyed, and Xing-Dong Yang. 2020. Capacitivo: Contact-based object recognition on interactive fabrics using capacitive sensing. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. 649–661.

[63] Yu-Chi Wu, Pei-Fan Chen, Zhi-Huang Hu, Chao-Hsu Chang, Gwo-Chuan Lee, and Wen-Ching Yu. 2009. A mobile health monitoring system using RFID ring-type pulse sensor. In *2009 Eighth IEEE International Conference on Dependable, Autonomic and Secure Computing*.

IEEE, 317–322.

[64] Dong Xu, Quan Yang, Feng Dong, and Sridhar Krishnaswamy. 2018. Evaluation of surface roughness of a machined metal surface based on laser speckle pattern. *The Journal of Engineering* 2018, 9 (2018), 773–778.

[65] Zihan Yan, Yuxiaotong Lin, Guanyun Wang, Yu Cai, Peng Cao, Haipeng Mi, and Yang Zhang. 2023. LaserShoes: Low-Cost Ground Surface Detection Using Laser Speckle Imaging. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.

[66] Xing-Dong Yang, Tovi Grossman, Daniel Wigdor, and George Fitzmaurice. 2012. Magic finger: always-available input through finger instrumentation. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*. 147–156.

[67] Hui-Shyong Yeo, Gergely Flamich, Patrick Schrempf, David Harris-Birtill, and Aaron Quigley. 2016. Radarcat: Radar categorization for input & interaction. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 833–841.

[68] Hui-Shyong Yeo, Juyoung Lee, Andrea Bianchi, David Harris-Birtill, and Aaron Quigley. 2017. Specam: Sensing surface color and material with the front-facing camera of a mobile device. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services*. 1–9.

[69] Hui-Shyong Yeo, Juyoung Lee, Hyung-il Kim, Aakar Gupta, Andrea Bianchi, Daniel Vogel, Hideki Koike, Woontack Woo, and Aaron Quigley. 2019. WRIST: Watch-Ring Interaction and Sensing Technique for Wrist Gestures and Macro-Micro Pointing. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*. 1–15.

[70] Cheng Zhang, Anandghan Waghmare, Pranav Kundra, Yiming Pu, Scott Gilliland, Thomas Ploetz, Thad E Starner, Omer T Inan, and Gregory D Abowd. 2017. FingerSound: Recognizing unistroke thumb gestures using a ring. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 1–19.

[71] Cheng Zhang, Xiaoxuan Wang, Anandghan Waghmare, Sumeet Jain, Thomas Ploetz, Omer T Inan, Thad E Starner, and Gregory D Abowd. 2017. FingOrbits: interaction with wearables using synchronized thumb movements. In *Proceedings of the 2017 ACM International Symposium on Wearable Computers*. 62–65.

[72] Cheng Zhang, Qiuyue Xue, Anandghan Waghmare, Ruichen Meng, Sumeet Jain, Yizeng Han, Xinyu Li, Kenneth Cunefare, Thomas Ploetz, Thad Starner, et al. 2018. FingerPing: Recognizing fine-grained hand poses using active acoustic on-body sensing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–10.

[73] Tengxiang Zhang, Xin Zeng, Yinshuai Zhang, Ke Sun, Yuntao Wang, and Yiqiang Chen. 2020. Thermalring: Gesture and tag inputs enabled by a thermal imaging smart ring. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.

[74] Xin Zhang, Karteek Kadimisetty, Kun Yin, Carlos Ruiz, Michael G Mauk, and Changchun Liu. 2019. Smart ring: a wearable device for hand hygiene compliance monitoring at the point-of-need. *Microsystem Technologies* 25, 8 (2019), 3105–3110.

[75] Yang Zhang, Gierad Laput, and Chris Harrison. 2018. Vibrosight: Long-range vibrometry for smart environment sensing. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 225–236.

[76] Yang Zhang, Sven Mayer, Jesse T Gonzalez, and Chris Harrison. 2021. Vibrosight++: City-scale sensing using existing retroreflective signs and markers. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14.

[77] Jan Zizka, Alex Olwal, and Ramesh Raskar. 2011. 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*. 489–498.