SpecTrans: Versatile Material Classification for Interaction with Textureless, Specular and Transparent Surfaces

Munehiko Sato[1]*, Shigeo Yoshida[2]*, Alex Olwal[1]†, Boxin Shi[1]‡,
Atsushi Hiyama[2], Tomohiro Tanikawa[2], Michitaka Hirose[2], Ramesh Raskar[1]
[1]MIT Media Lab, Cambridge, MA, USA
[2]The University of Tokyo, Bunkyo, Tokyo, Japan
{munehiko, olwal, shiboxin, raskar}@media.mit.edu
{shigeodayo, atsushi, tani, hirose}@cyber.t.u-tokyo.ac.jp

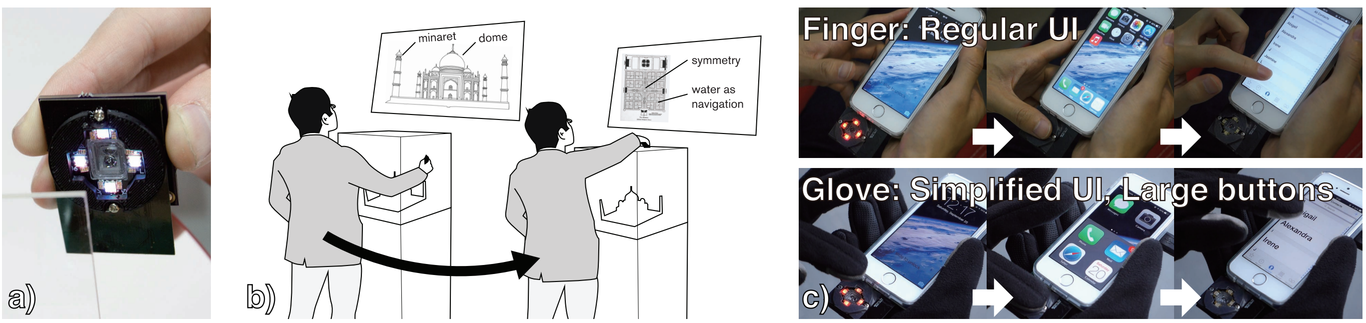

Figure 1: SpecTrans is a versatile material classification technique that enables various applications with textureless, specular and transparent surfaces. a) SpecTrans sensor in action. b) A transparent showcase made of six different materials that shows different display contents according to which side the user is touching. c) A context-aware smartphone home button that adapts the appearance and behavior of the user interface. It shows bigger icons when a user touches the button with a winter glove made from conductive-threads.

ABSTRACT

Surface and object recognition is of significant importance in ubiquitous and wearable computing. While various techniques exist to infer context from material properties and appearance, they are typically neither designed for real-time applications nor for optically complex surfaces that may be specular, textureless, and even transparent. These materials are, however, becoming increasingly relevant in HCI for transparent displays, interactive surfaces, and ubiquitous computing.

We present SpecTrans, a new sensing technology for surface classification of exotic materials, such as glass, transparent plastic, and metal. The proposed technique extracts optical features by employing laser and multi-directional, multi-spectral LED illumination that leverages the material's optical properties. The sensor hardware is small in size, and the proposed classification method requires significantly lower computational cost than conventional image-based methods, which use texture features or reflectance analysis, thereby providing real-time performance for ubiquitous computing.

Our evaluation of the sensing technique for nine different transparent materials, including air, shows a promising recognition rate of 99.0%. We demonstrate a variety of possible applications using SpecTrans' capabilities.

Categories and Subject Descriptors: H.5.2 User Interfaces: Input devices and strategies

Keywords: Sensors; material classification; laser speckle; multi-spectral sensing; ubiquitous computing; context-aware mobile computing

*The first two authors contributed equally to this work.
†Secondary affiliations: KTH – Royal Institute of Technology (Stockholm, Sweden); Google [x] (Mountain View, CA, USA)
‡Secondary affiliations: Singapore University of Technology and Design (Singapore)

CHI 2015, April 18–23, 2015, Seoul, Republic of Korea.

ACM 978-1-4503-3145-6/15/04 ...$15.00.
http://dx.doi.org/10.1145/2702123.2702169

INTRODUCTION

Context is extremely important in Human-Computer Interaction (HCI). It is particularly relevant for interactions in the real world, such as for ubiquitous computing, wearable devices, organic user interfaces and mixed reality. There are many different sensing techniques that enable an interface to adapt its appearance, behavior, functionality and representation to best fit the user's context. Hinckley et al. [8], for example, describe how a mobile device can adapt its interface to the user based on proximity, touch, grip and tilt.

Other works have focused on inferring information from the interaction and input surfaces themselves. Harrison and Hudson [7], for example, use a single photoresistor with multi-spectral illumination to allow a mobile device to identify the surfaces on which it is placed.

Magic Finger [35], on the other hand, uses a NanEye camera to classify ordinary materials by their textures. Similar to the Anoto technology[1] that uses printed microdot patterns, Magic Finger makes it possible to recognize small letters (2 pt) and fiducial markers printed by an office laser printer to trigger functionality via the encoded data.

Emerging interactive surfaces, transparent displays, smart homes, and other real-world environments, introduce new challenges for this type of sensing, which cannot be addressed by the above-mentioned techniques. It is often not feasible to modify surfaces, such as store display windows, projection screens, or interactive tabletops, for practical, technical or aesthetic reasons. We may, for example, wish to maintain a see-through display's transparency while being able to seamlessly support detection and recognition. Many popular materials today, such as machined aluminum and plastics, not only lack visual texture, but may also be highly specular or transparent.

To support context-awareness in these scenarios, we introduce SpecTrans, a new sensing technique that can capture rich information about the material's optical properties at high speed with minimal computation. SpecTrans combines a fast image sensor with multi-directional, multi-spectral imaging to capture the varying optical properties of different materials under different lighting conditions (Figure 2). This makes it possible to extract simple, yet efficient, features for material classification that cannot be observed given a single image under a fixed illumination. Our sensing hardware contains an optical mouse image sensor, a laser emitter, and 20 light emitting diodes (LEDs), clustered with five different wavelengths in four directions. The image sensor's on-chip image processing unit can execute several fundamental operations before transmitting the data to a microcontroller for further processing, which enables real-time feature extraction.

CONTRIBUTIONS

This paper proposes a novel method for classifying surfaces of different materials in a mobile context, using a small and fast image sensor with laser and multi-directional, multi-spectral LED illumination:

1. We introduce a novel material classification technique based on the laser speckle and multi-directional, multi-spectral capture of optical material properties.
2. We describe the design and implementation of sensor hardware and software that enable embedded devices with real-time material classification with micro-controllers.

[1] `http://www.anoto.com/`

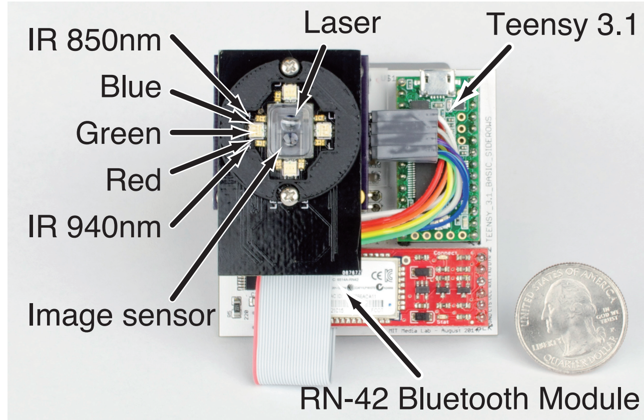

Figure 2: SpecTrans sensor (bottom view). Four LED clusters with five different wavelengths (red, green, blue, and two infra-red (IR) wavelengths) are arranged in different directions around the sensor. Together with the laser, SpecTrans uses 21 light sources in different configurations and sequences to capture a rich set of features at high speeds.

3. We report the results from classifying 21 common materials, (including nine transparent ones), and show how our SpecTrans is more versatile than image-based techniques.
4. We describe application scenarios and interactions that are enabled by the unique capabilities of real-time, embedded material sensing.

RELATED WORK AND APPROACHES

Material classification and context-aware computing have significant importance to HCI. Numerous approaches have been proposed to enable computers, devices and robots to understand the physical world so they can support responsive and interactive systems.

Material Classification Using Imaging and Lighting

For a specific lighting and viewpoint, the visual appearance of an opaque object, with a known shape, is determined by its reflectance properties, which are described by the 4D Bidirectional Reflectance Distribution Function (BRDF). The BRDF is highly related to surface material properties, and material recognition is more easily tractable if the BRDF is known. However, from a single image with an unknown object shape and illumination, this task is extremely challenging, since capture of the complete BRDF requires knowledge of object shape and images from all lighting directions and viewpoints [17]. When an object is translucent or transparent, its reflectance becomes even more complicated and needs to be described by a higher dimensional model [11].

Material classification can be difficult using a single natural image [27], but in a controlled setup with a specially designed imaging device, it becomes much more tractable even with only color and texture information. A miniature RGB camera could, for example, be worn on the fingertip to sense environmental and artificial textures, such as printed, microscopic characters or fiducials, to enable context-triggered taps and gestural interactions [35]. Similar image-based surface classification techniques are used in robotic hand manipulation tasks using a laser optical mouse sensor [16]. However, these methods are designed for surfaces with distinctive textures, and do not work on transparent surfaces.

There are existing works that explore the use of multi-directional lighting, in addition to color and texture, to estimate microscale 3D surface structure for improved material classification [15]. Using multi-spectral and multi-directional illumination to create discriminative illumination conditions enables the classification of complex materials, such as unpainted raw metals [4]. Our approach shares the same theoretical support as that of [4], but we use a simplified setup in terms of diversities of spectra and lighting directions. By carefully selecting representative spectra and lighting directions, we focus on capturing enough features for real-time determination of material categories for HCI applications, instead of producing a pixel-wise classification.

Our work emphasizes real-time material classification for mobile contexts and is, therefore, different from BRDF measurement devices, such as [1] and [25], which are designed for accurate BRDF capture. The former device is a 0.5 m hemisphere, which is neither portable nor real-time, while the latter requires a pre-defined BRDF chart and 30 seconds for data capture. Our applications do not require capture of complete BRDF information since we focus on the measurement of sufficient information for real-time classification of challenging materials.

There are more sophisticated material and content detection techniques for scientific and industrial applications. For example, millimeter wave and terahertz technology are being used for scientific and security purposes to detect materials from a distance [13]. These sensing methods are, however, less applicable to wearable and ubiquitous computing given their complexity, cost, size, and power requirements.

The closest work to ours is [7], which shows how a set of non-directional ultraviolet, red, green, blue and infrared LEDs, in combination with a photoresistor and a light sensor IC can measure simplified multi-spectral characteristics of object surfaces to enable context-aware mobile devices. Our method captures more complex surface features in considerably more conditions, beyond simply texture and color, using 104 different lighting conditions, including coherent laser illumination. This enables interaction with not only more complex materials, but also specular and transparent surfaces in particular, where layering becomes an additional possibility.

Environmentally Embedded Markers and Tags

Markers and tags are often used to enable augmented reality (AR) and context-aware ubiquitous computing. Image-based fiducial markers, e.g., bar codes, QR codes, and ARToolKit tags [12], are popular for camera-based AR. The Anoto technology is based on a pen with an embedded small-size video camera, to track the absolute position of its tip on a paper with printed microdot patterns. Encoded Reality[2] proposes encoding patterns in physical objects using laser etching and other fabrication techniques. These approaches require special objects or surface modifications to enable recognition. Bokode [19] uses tiny optical tags that can be read using conventional cameras from a distance, whereas other work uses small active and passive tags that communicate using radio frequencies [18], time-coded projection patterns [23], or spatially-coded projection patterns [34]. Additive and subtractive manufacturing technologies, e.g., 3D printing, enable embedding custom 3D markers inside objects that can be decoded with a terahertz scanner [33], or to create optical components such as light-guide pipes and optical sensors [32].

The need for special tags in the environment, however, limits possible applications and scalability through the need of special infrastructures and instrumentation.

Context Classification for Mobile and Wearable Scenarios

There is a tremendous body of work on user and mobile device context detection. Personal information devices, such as smartphones, are intimate objects that can infer rich information about user context. The device can use prediction models for the actions users are about to perform, based on how the device is grasped or held, for example, through the use of capacitive sensor matrices [30]. Ichikawa et al. surveyed where people keep their phones [10], while Kunze et al. investigated device location detection using accelerometers [14]. More recently, Wiese et al. demonstrated sensor fusion of multiple sensing dimensions, using a capacitive sensor matrix, multi-spectral light sensor, and an accelerometer [31].

Interactions with Transparent Materials and Surfaces Using Optical Properties

Numerous interaction techniques have been proposed to interact with transparent surfaces and objects, utilizing their optical properties for sensing and display. Sato et al. exploited the photoelastic effect from transparent elastic materials when they are deformed to enable freeform tangible interaction on a display [26]. The use of index-matched liquid makes an elastic container with glass particles optically transparent, to enable both shape-sensing and projection through the volume, in a free-form jammable user interface [3]. Optical fibers and light guide pipes are classical optics components, but they can be used in various applications, such as capturing fingerprints on a touchscreen [9], fabricating interactive 3D printed objects with curved display surfaces [2, 21] or for sensing [32]. Rekimoto et al. [24] explore transparent tiles for tangible interaction with embedded affordances and position sensing.

We note that the sensing of optical properties for transparent materials has interesting potential for spatial and mobile platforms where issues like size, power, computational cost and complexity must be taken into consideration.

[2] `http://viral.media.mit.edu/projects/encoded_reality/`

Laser Speckle Sensing for Interaction
Previous work has also explored speckle sensing for user interfaces. SpeckleSense [36] explores the use of high-speed motion sensing using laser speckle for input devices, gestural interaction and spatially aware mobile devices. Olwal et al. [20] described a platform that uses speckle sensing for gestural interaction with embedded electronics. Davis et al. [28] employed speckle to detect surface tampering.

SPECTRANS: DESIGN OF MULTI-SPECTRAL, MULTI-DIRECTIONAL MATERIAL CLASSIFICATION
In SpecTrans, our goal is to overcome the limitations of camera-based capture, with an embedded sensor that can detect surface material using a low-cost laser mouse sensor and multi-directional, multi-spectral illumination. We leverage on-chip image processing to support real-time material classification from a relatively small, but comprehensive, set of features that can be processed on embedded platforms.

Surface Sensing with Coherent and Incoherent Light
Three major techniques are used in commercial optical mice today. Laser Doppler Interferometry [22] relies on the interference of laser light, and tracks velocity without an image sensor. This hardware is, therefore, less suitable for capturing the type of surface features we are interested in. In contrast, LED-illuminated and laser-illuminated mice use small image sensors (matrix of photo sensors) on the chip. The light source is also embedded with lenses. We combine these two lighting techniques with multi-directional lighting for robustness and high surface recognition rates.

Surface Feature Capture
LED-based optical mice illuminate the operating surface, e.g., a mouse pad or table, with collimated incoherent light. The surface image is captured by the sensor, and 2D displacement between two consecutive frames is calculated by an on-chip processor at very high speeds (1–10 kHz). A collimating lens after the LED illumination helps capture surface features.

Laser Speckle Capture
Laser-based optical mice use coherent light to achieve stable tracking on various surfaces, including very shiny and flat surfaces. A laser diode on a sensor chip emits a coherent laser beam, which travels through a collimating lens and illuminates the surface. The reflected light from the tracking surface creates a fine *speckle pattern*, which the image sensor tracks, instead of the surface texture. The speckle pattern results from interference of the reflected wavefronts due to microscopic roughness on the surface that varies their path lengths [36]; therefore, it can exploit small variances in material surfaces that are much smaller than the pixel pitch. We note that the speckle pattern is a statistic of phenomena at the image sensor plane.

We take advantage of this speckle phenomena to capture features of *textureless* surfaces, such as transparent plastic sheets, glass sheets, and shiny metal.

Multi-directional and Multi-spectral Lighting
We use multi-directional and multi-spectral lighting to acquire more surface detail. Leveraging the fast mouse sensor enables us to quickly take multiple samples of surface reflection.

Highly specular surfaces, e.g., metal, reflect incoming light like a mirror, while diffused surfaces reflect light more uniformly. Transparent surfaces interact with light in a complicated way, including both reflection and refraction. Inspired by [4], which indicated that multi-directional and multi-spectral information can provide discriminative features for textureless material classification, we design our sensor with multi-spectral LEDs at different locations.

Multi-directional illumination is key in the design of our sensor. Transparent and highly specular materials have complicated, but discriminative, BRDFs concentered around the center of the specular lobe. These delicate differences can only be captured with carefully designed multi-directional lighting and a light sensor array, rather than a single pixel as in [7]. SpecTrans enables such sampling around the center of the specular lobe by using an image sensor, under various lighting directions and spectrums. The 2D intensity map of all pixels (texture features) is condensed to three key BRDF parameters with no computational cost for a microcontroller. Our method is also fundamentally different from [35] that relies on texture patterns from the image sensor, which add computational and bandwidth constraints.

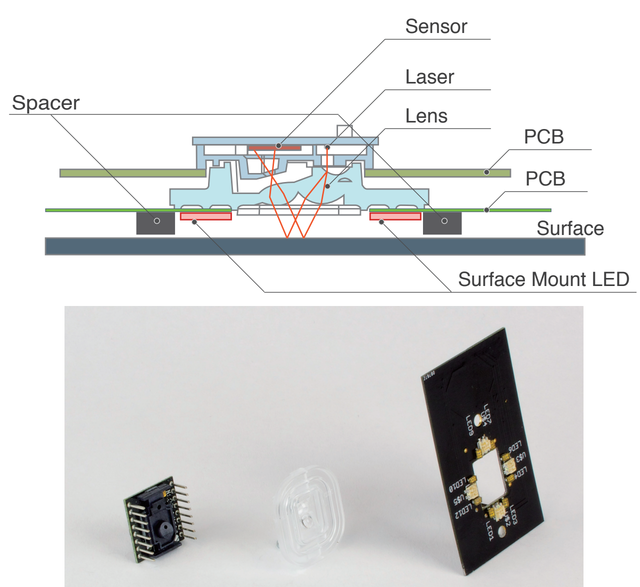

Figure 3: Sectional and exploded view of the proposed SpecTrans sensor. Additional 20 LEDs in 5 different wavelengths and from 4 directions are arranged on a custom PCB.

Sensing System Design
Hardware Design
We use a high-speed laser optical mouse sensor (Avago ADNS-9500[3]) with a small form factor lens (Avago ADNS-6190-002). The lens is a molded single-piece plastic with two optical lens components — one for collimating a laser

[3] `http://www.pixart.com.tw/`

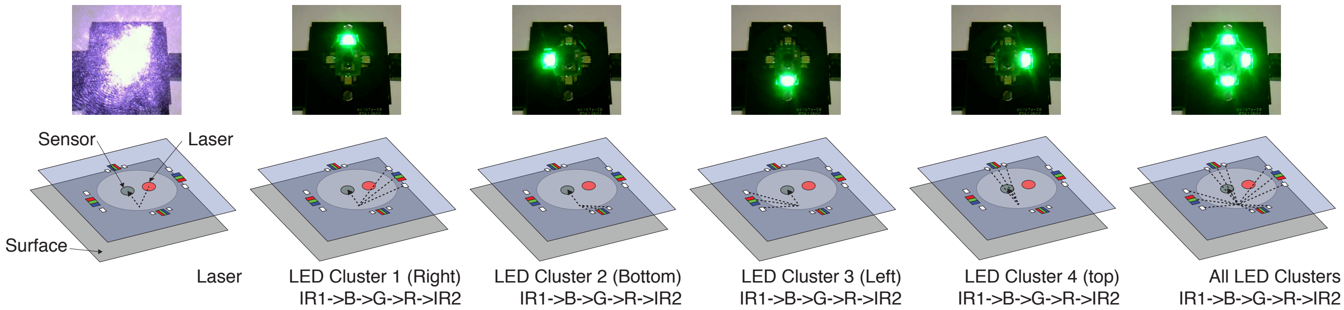

Figure 4: Different lighting configurations. Four clusters of five LEDs are located around the image sensor and laser emitter.

beam (Vertical-cavity surface-emitting laser (VCSEL) with 832-865 nm peak wavelength) to illuminate the tracking surface, and the other for imaging the tracking surface in the 30×30 pixel image sensor. We extend these optical elements with four LED clusters placed at different positions, each containing LEDs at different peak wave lengths of blue (475 nm), green (525 nm), red (621 nm), and two infrared (850 and 940 nm) for a total of 20 LEDs (see Figure 2 and 3 [4]). We chose health-friendly (non-ultraviolet) and sensitive wavelengths for the sensor. The LED clusters are arranged in a square to enable multi-directional lighting from various angles while keeping the compactness of the sensor module.

We created two versions of our control boards using Arduino-compatible 32-bit ARM microcontrollers: a) Teensy 3.1 (Cortex-M4, 96MHz) with Microchip RN-42 Bluetooth module, and b) Spark Core (Cortex-M3, 72MHz) with on-board WiFi module. Our prototype sensing system is completely wireless and its dimension is roughly $60\times60\times40$ mm, including a rechargeable battery. The optical part of the system (sensor chip, lenses, and LEDs) is $20\times18\times11$ mm. It enables the embedding of the sensor in handheld or wearable devices, as well as, in the environment.

We set the LED driving current to 30 mA using constant current regulator ICs. The power consumption of the LEDs can be minimized by setting the same times for LED illumination and sensor exposure time. The depth of field of the ADNS-9500 sensor is designed to be very small, i.e., +/- 0.22 mm as the typical distance from the lens reference plane to the tracking surface. We use a thin (0.6 mm) PCB to give maximum distance between surfaces and the illumination. A black plastic spacer keeps the sensor-surface distance constant (see Figure 2 and Figure 3 top.)

Different Lighting and Shutter Speed

As previously described, we use both the collimated coherent laser and incoherent LED for classification. Figure 4 illustrates the different lighting patterns that we employ. As we use LEDs with five different wavelengths, the total number of lighting patterns are 26 (1 laser + 5 wavelengths) $\times$ (4 LED clusters + all LED clusters).

[4]This diagram was created based on **`http://www.electronicproducts.com/Sensors_and_Transducers/Image_Sensors_and_Optical_Detectors/Optical_mouse_technology_Here_to_stay_still_evolving.aspx`** and the datasheet of Avago ADNS-9500 sensor.

While the chip normally uses auto exposure for mouse applications, our objective is to capture multiple exposures of the surface features under controlled conditions. We, therefore, configured the sensor to use manual shutter speeds in order to obtain high dynamic range images to use in the classification. We chose four fixed exposure periods: 0.83, 6.64, 53.12, and 424.96 μs. These periods correspond to 39, 312, 2496, and 19968 clock periods of the sensor chip's internal oscillator at 47MHz. Thanks to this multi-exposure setup, we can capture a high dynamic range of surface reflections from dark carpet to shiny metal, and black synthetic fur to transparent glass. Our current implementation can capture four features at four different exposure setups with 26 lighting conditions (416 features in total) in merely 117 ms.

Sensing Capabilities and Surface Features

The speckle pattern is a random image with low information density, which makes it possible to reduce the captured image to the statistical parameters that describe it.

Similar to the sensors used in previous work (e.g., [20]), the sensor we use has various on-chip-computation features, such as displacement $(\Delta x, \Delta y)$, mouse lift detection, darkest/brightest/average pixel intensity, surface quality, and raw pixel values. Surface quality is a measure of valid feature points captured by the ADNS-9500 sensor. It provides a useful index for roughness of the surface. We experimentally determined that pixel intensity statistics and surface quality are the best descriptors of material features.

SpecTrans exploits on-board processing to extract three key parameters and reduces the amount of data to transfer from 900 to four bytes:

- Average pixel intensity $\rightarrow$ Average brightness
- Darkest and brightest pixel $\rightarrow$ Contrast
- Surface quality $\rightarrow$ Pattern granularity

Transparent or highly specular surfaces provide almost no texture information for the sensor to capture under collimated LED illumination. Our proposed multi-directional illumination overcomes this challenge by capturing complicated and discriminative material features of such surfaces.

Our embedded multi-spectral, multi-directional sampling relies on the sensor's high-speed capture, enabled by the on-chip computation of these features. This makes these features available to microcontrollers at a low bandwidth, since there

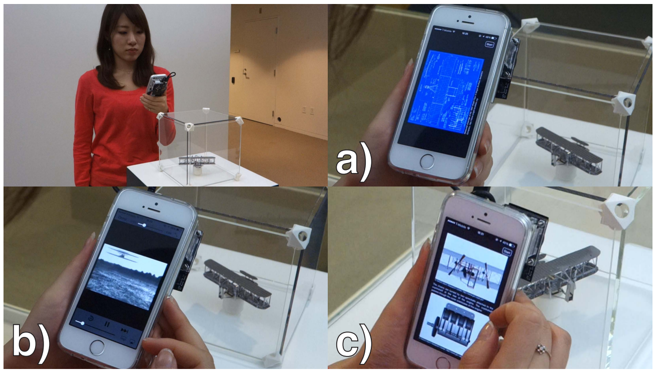

Figure 5: Three different contents are associated with each side of the showcase. a) Top-view blueprint of the plane. b) Video footage of Wright Flyer flying toward the video camera. c) Animation showing how the engine and propeller work.

is no image transfer or processing outside the sensor chip. By choosing these features, we ensure scalability, since their suitability for on-chip computation makes them available in most types of optical mouse sensors.

APPLICATIONS AND INTERACTIONS

SpecTrans is unique in that it can classify transparent and specular surfaces. Its minimal computational cost and real-time performance make it suitable for many different types of embedded devices and scenarios, some of which we will discuss below. All of the described demo applications were implemented and work in real time.

The intent of this paper is to show some possibilities with SpecTrans. It is a complementary technology that reduces the need for switches, buttons and other active electronics, but it is not intended to replace them entirely. SpecTrans is designed for location awareness, but it also enables the embedding of invisible information in raw materials without extra cost, electronics or infrastructure.

Architects, industrial designers, interior designers, and digital media artists continually seek new tools to realize their digital creations in existing materials. For example, museum dioramas are designed to reproduce wildlife scenarios; however, QR codes, AR markers, RFID/NFC tags, and caption boards risk interfering with the exhibit and its artifacts. There are similar challenges for store display windows and price tags in jewelry stores. SpecTrans expands the interaction vocabulary for such scenarios.

Encoding Imperceptible IDs within Transparent Surfaces

Interacting with transparent surfaces and displays enables extremely useful applications. By sensing imperceptible optical characteristics in the surfaces, the system maintains its transparent characteristics, while being able to respond to different user interactions.

Transparent Showcase: Adaptive Presentation

The six sides of our transparent showcase use different transparent materials with imperceptible properties. Our system can accurately distinguish between the six transparent surfaces of the cube, based on real-time classification results. Using this information, we map different types of content and interactions to the individual faces of the showcase, according to what has been sensed by our device (Figure 1 b) and Figure 5.)

A smartphone with our embedded sensor displays unique contents based on which side it is held against. This example shows a model of a Wright Flyer in the showcase (Figure 5.) By scanning the front face of the showcase, the screen displays old footage of the plane flying towards the user. When the user scans the right face of the showcase, where the propeller of the plane is most visible, an animation shows how the engine and propeller work. Finally, when the user touches the top of the showcase, a top-view blueprint of the airplane appears on the screen. While it is unnoticeable to the user, each transparent surface naturally encodes an ID in the material's optical properties.

Transparent Control Pads: Changing Control and Viewports

We have also explored scenarios that combine SpecTrans with tracking on transparent surfaces. Figure 6 shows how the user can physically switch a mouse pad to change the mapping of the mouse's movement between zooming, panning and rotation. Figure 7 shows how different segments of a larger mouse pad can be mapped to viewports or virtual desktops to leverage a user's proprioception and spatial memory. In this example, placing the mouse in a segment activates either an overview or one of three perspectives.

Transparent Overlays: Tangible Data Filters

Figure 8 shows how a physical overlay can be detected by SpecTrans to reveal additional information and personalize the UI on a tablet, without interfering with the display. This approach does not require an alignment of tags and sensors, as for NFC, and fully maintains the transparent properties of the passive material. The drawback is, of course, that anyone with access to the material can replicate the filter, but such security issues are beyond the scope of this paper.

Context-Aware Input: Adapting UI to Skin or Glove

Certain input devices, like optical trackballs (e.g., BlackBerry Curve 8520), fingerprint readers (common on many laptops), and fingerprint-sensing buttons (e.g., Apple TouchID) rely on the user's contact with a sensor. We augmented an iPhone 5S with SpecTrans to infer which material is placed on top of it. This information can be used to adapt the appearance and behavior of the UI for subsequent interactions. The UI could, for example, make the UI elements larger and adjust touch-screen sensitivity if it detected that the user was wearing gloves by sensing leather, wool, or cotton instead of skin. Figure 1 c) shows how our current prototype changes the UI to a few large buttons for apps that are relevant outdoors, when gloves are detected.

TECHNICAL EVALUATION

SpecTrans extends previous research that used color- and texture-based sensing for context-aware interaction [7, 35] by extracting invisible information directly from the environment, i.e., transparent raw materials. We evaluate the surface

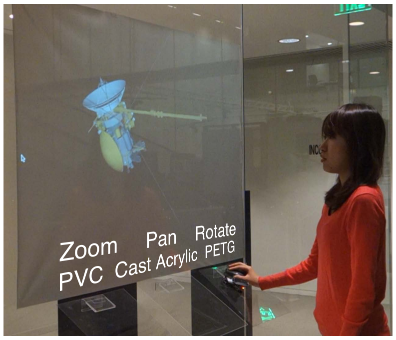

Figure 6: Controller for a floating transparent screen. Different controls (Pan, Rotate, and Zoom) are associated with a set of mouse pads with different transparent materials.

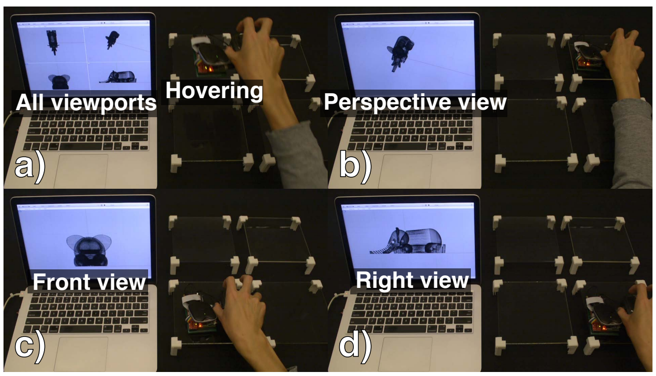

Figure 7: Four viewports of Rhinoceros 3D CAD software are associated with four mouse pads with four different materials. 1) Showing all viewports when hovering the mouse in the air. 2) Perspective view. 3) Front view. 4) Right view.

classification accuracy of our sensor with three types of material for quantitative comparison: transparent material, synthetic animal fur and skin, and surfaces typically found in home and office environments.

We report on the classification accuracy and investigate how each lighting condition contributes to the classification. We selected the following materials for each test scenario:

Nine transparent materials: Polyethylene terephthalate glycol-modified (PETG), polyvinyl chloride (PVC), cast acrylic, cellulose acetate, extruded acrylic, glass, microscope slide, polycarbonate, and air (no material). (See Figure 9.)

Five synthetic animal furs and skins from a picture book: Lion fur, panda fur, iridescent colored lizard skin, elephant skin, and zebra fur.

Seven office and living spaces: Office desk, aluminum laptop palm rest, chair seat (fabric), chair arm rest (rubber), fabric mouse pad, plastic mouse pad, and carpet.

Data capture
Data capture was done by placing the sensor on top of the surface sample. For each sample, we captured 20 data instances

Figure 8: Changing the behaviors of a tablet device by overlaying transparent sheets made from different materials. It shows numerical data for the number of items in stock, breakdown costs, profit ratio, and remarks only for sales personnel, while a customer can only see the sum.

while moving the sensor gently over the surface. This process was repeated five times for all of the materials. In total, we obtained 100 samples for each material.

For transparent materials, we add "Air" (a surface without any material) to compare with other transparent plastics and glasses. All experiments were conducted in an ordinary office environment with fluorescent ambient lighting. The transparent surfaces are cleaned with lens cleaning wipes before being captured.

Classification
We used the Support Vector Machine (SVM) classifier by Sequential Minimal Optimization (SMO), implemented in the Weka Toolkit [5] to perform classification. As features, we used the four parameters (surface quality, average pixel intensity, darkest and brightest pixel intensity) directly read from the sensor for each lighting configuration and shutter speed configuration. In total, we had 416 features (4 parameters $\times$ 4 exposure times $\times$ 26 lighting patterns) for one data point.

Performance Evaluation
We performed a ten-fold cross validation for three scenarios independently to evaluate how accurately we could classify the materials. The average correct classification rate is shown in Figure 10 with ten different lighting conditions in Table 1 to prove the necessity of introducing multi-spectral and multi-directional lighting.

We achieved a high recognition rate (above 99.0%) on average for all of the materials with our comprehensive setup **J** (SpecTrans), when all lighting directions and spectra were considered. Setup **C** was designed to mimic a simple multi-spectrum LED method. The general trend we observed is that greater variance in lighting directions and more spectral configurations will produce higher classification accuracy. For transparent materials in Figure 10 a), the simple non-directional multi-spectrum LED sensing (**C**) showed a low correct classification rate of 65.44%, which verifies the necessity of our comprehensive setup. By adding laser to the illumination, correct classification rates increase under all

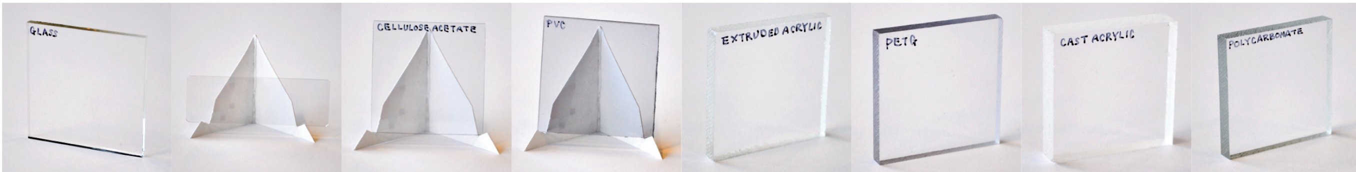

Figure 9: Close-up views of eight transparent materials. All pictures were taken with the same camera setting and lighting condition (Nikon D5000, F 5.6, exposure 1/100 in an ordinary office environment.)

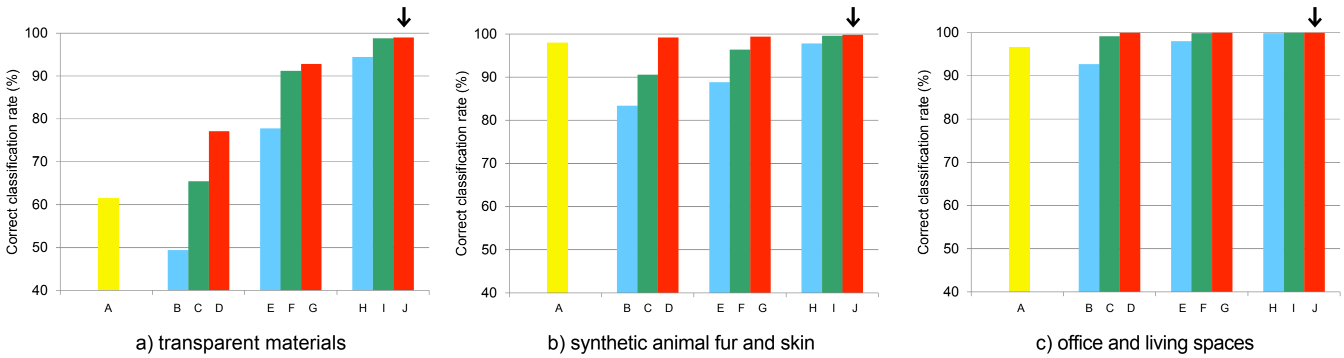

Figure 10: Correct classification rate: a) transparent materials, b) synthetic animal fur and skin, and c) office and living spaces. For each scenario, ten lighting configurations in Table 1 were tested. **J** (with an arrow) is the SpecTrans.

conditions as seen in Figure 10 (green and red bars). With transparent materials (Figure 10 a)), the correct classification rate increased by adding laser from 98.78% (**I**) to 99.0% (**J**), and from 65.44% (**C**) to 77.11% (**D**). Also, Laser only (**A**) surpassed single LED (**B**).

Table 1: Lighting conditions for performance evaluation.

Laser Only	A

Spectrum \ Location	LED Cluster # 1	LED Cluster # 1 and 2	LED Cluster # 1, 2, 3, 4, and All
IR-850 nm	**B**	**E**	**H**
IR-850 nm, IR-940 nm, R, G, and B	**C**	**F**	**I**
IR-850 nm, IR-940 nm, R, G, and B, and Laser	**D**	**G**	**J** (SpecTrans)

Figure 11 shows the confusion matrices for the transparent materials under three (out of ten) different lighting configurations. These matrices show how the combination of laser and multi-directional and multi-spectral illuminations contributes to more accurate classifications.

LIMITATIONS AND FUTURE WORK

Our current implementation used a dedicated optical sensor for material classification. It would, however, be possible to temporally multiplex material classification with motion sensing using a single optical sensor and chip. This could, for example, be done by inferring a change of context, such as detecting lift-off or landing on surfaces using the sensor chip's lift detection function or inertial sensors. Thus, after detecting lift-off, a mouse could switch into the SpecTrans material classification mode upon surface contact and after identification returns to the tracking mode again. Another approach is to alternate between classification and tracking, making use of high-speed tracking and material classification.

SpecTrans requires additional illumination for multi-spectral, multi-directional captures. For this project, we used 20 LEDs and a laser in our general SpecTrans device; however, fewer LEDs could be used for less challenging materials.

We currently can capture all 11 features under 104 different lighting conditions and exposures, which are transmitted to the micro-controller for further processing. The capture process takes 117 ms, which is fairly fast and sufficient for real-time applications. In future application scenarios, we anticipate that the capture will use fewer LEDs and exposure settings, optimized for the specific materials of interest. Our current implementation, with a very wide dynamic range, was designed to be versatile from clear glass to matte black paint. Single exposure is sufficient for specific applications, such as UI with transparent surfaces, resulting in a four times faster data capture. This, in combination with a faster microcontroller and optimized data transfer, will yield significantly higher speeds. Another approach is to use multiple sensors that can collect different features in parallel.

For certain applications it may, however, be interesting to expand the illumination for wider, multi-spectral capture using both additional LEDs and lasers. We note that the VCSEL laser has a vertical configuration on the silicon, which makes it an ideal building block for dense array arrangements with multi-spectral lasers. For applications that require a complete

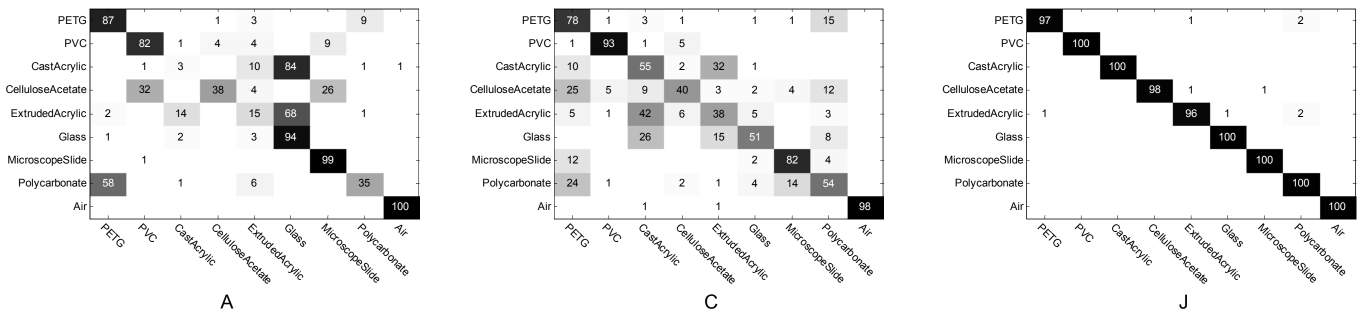

Figure 11: From left to right, confusion matrices for lighting condition **A**, **C**, and **J** with transparent materials (Figure 10 a)). The value is the number of correctly classified examples among 100 samples. We omitted the zero values in the matrices for better clarity.

BRDF capture, the system could be coupled with more advanced optical elements (e.g., [6]).

We are also interested in applying our sensor to other applications, such as wearable input for health purposes, e.g., for detecting body location or capturing skin subsurface imaging. We are, in particular, interested in using SpecTrans for non-invasive health screenings with techniques such as laser speckle imaging and laser doppler flowmetry to assess microcirculatory functions and blood flow [29].

The SpecTrans multispectral configuration also addresses challenges with ambient light, which are not an issue in prior work that focuses on opaque surfaces. While environmental lights, such as sunlight, incandescent, fluorescent, and LED light sources, might affect the sensor, their incoherence means that they do not interfere with the laser speckle pattern. Additionally, in future work, ambient light could be filtered out using conventional modulation noise reduction methods, which are widely used, for example, in TV remote controls.

CONCLUSIONS

In this paper, we presented SpecTrans, a new sensing platform to provide context-aware user interfaces using rapid, multi-spectral, multi-directional capture and classification of optically challenging materials. Our technology can be used alone or in conjunction with other sensing approaches to enable rich and robust context awareness for ubiquitous, mobile and wearable computing. SpecTrans uses a low-resolution, ultra-fast, image sensor to capture a small set of optical features under varying illumination conditions. Hardware-accelerated image processing calculates the features on the image sensor's chip, allowing it to be implemented as a low-power, real-time sensing platform with a micro-controller.

SpecTrans can bring context-awareness to numerous exotic surfaces that currently have to rely on markers, patterns or other modifications. The versatile sensing makes it uniquely possible to classify materials, such as acrylic, polycarbonate, glass, and metal, which is particularly relevant for emerging exotic display and interaction surfaces. SpecTrans can, however, also be used for other materials with simpler optical characteristics. Our evaluation shows 99.0% accuracy for classifying a set of eight + one (air) surfaces, including a range of diffuse, as well as, highly reflecting materials. These promising results indicate the potential for SpecTrans as a fast, low-power, low-cost technology for mobile and ubiquitous context-awareness.

ACKNOWLEDGMENTS

The authors wish to thank Jessica Qian for her technical assistance, and Ryan Mather, John Klein, Xin Liu, and Yueping Wang for their valuable feedback on applications. Munehiko Sato is financially supported in part by the Microsoft Research collaborative research (CORE) program, the Nakayama Foundation for Human Science, and the Mizuho Foundation for the Promotion of Science. Shigeo Yoshida is funded by a JSPS research fellowship, Boxin Shi is supported by the SUTD-MIT joint postdoctoral program, and Alex Olwal is supported in part by the Swedish Research Council. Finally, we thank our reviewers for insightful feedback that helped us improve this paper.

REFERENCES

1. Ben-Ezra, M., Wang, J., and Wilburn, B. An LED-only BRDF measurement device. In *Proc. IEEE CVPR '08*, 1–8.

2. Brockmeyer, E., Poupyrev, I., and Hudson, S. Papillon: Designing curved display surfaces with printed optics. In *Proc. ACM UIST '13*, 457–462.

3. Follmer, S., Leithinger, D., Olwal, A., Cheng, N., and Ishii, H. Jamming User Interfaces: Programmable Particle Stiffness and Sensing for Malleable and Shape-Changing Devices. In *Proc. ACM UIST '12*, 519–528.

4. Gu, J., and Liu, C. Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral BRDF. In *Proc. IEEE CVPR '12*, 797–804.

5. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., and Witten, I. H. The WEKA data mining software: an update. *SIGKDD Explorations 11*, 1 (2009), 10–18.

6. Han, J. Y., and Perlin, K. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. on Graph. 22*, 3 (2003), 741–748.
7. Harrison, C., and Hudson, S. E. Lightweight material detection for placement-aware mobile computing. In *Proc. ACM UIST '08*, 279–282.
8. Hinckley, K., Pierce, J., Sinclair, M., and Horvitz, E. Sensing techniques for mobile interaction. In *Proc. ACM UIST '00*, 91–100.
9. Holz, C., and Baudisch, P. Fiberio: a touchscreen that senses fingerprints. In *Proc. ACM UIST '13*, 41–50.
10. Ichikawa, F., Chipchase, J., and Grignani, R. Where's the phone? A study of mobile phone location in public spaces. In *Proc. IEEE Mobility Conf. '05*, 142–142.
11. Jensen, H. W., Marschner, S. R., Levoy, M., and Hanrahan, P. A practical model for subsurface light transport. In *Proc. ACM SIGGRAPH '01*, 511–518.
12. Kato, H., and Billinghurst, M. Marker tracking and HMD calibration for a video-based augmented reality conferencing system. In *Proc. IWAR '99*, 85–94.
13. Kemp, M. C. Millimetre wave and terahertz technology for the detection of concealed threats: a review. In *Optics/Photonics in Security and Defence*, Intl. Society for Optics and Photonics (2006), 64020D–19.
14. Kunze, K., Lukowicz, P., Junker, H., and Tröster, G. Where am I: Recognizing On-Body Positions of Wearable Sensors. In *Proc. Intl. Conf. on Location- and Context-Awareness '05*, 264–275.
15. Li, R., and Adelson, E. H. Sensing and recognizing surface textures Using a GelSight sensor. In *Proc. IEEE CVPR '13*, 1241–1247.
16. Maldonado, A., Alvarez, H., and Beetz, M. Improving robot manipulation through fingertip perception. In *Proc. IEEE/RSJ IROS '12*, 2947–2954.
17. Matusik, W., Pfister, H., Brand, M., and McMillan, L. A data-driven reflectance model. *ACM Trans. on Graph. 22*, 3 (2003), 759–769.
18. Michahelles, F., Thiesse, F., Schmidt, A., and Williams, J. R. Pervasive RFID and Near Field Communication Technology. *IEEE Pervasive Computing 6*, 3 (2007), 94–96, c3.
19. Mohan, A., Woo, G., Hiura, S., Smithwick, Q., and Raskar, R. Bokode: Imperceptible Visual Tags for Camera Based Interaction from a Distance. *ACM Trans. on Graph. 28*, 3 (2009).
20. Olwal, A., Bardagjy, A., Zizka, J., and Raskar, R. Speckleeye: Gestural interaction for embedded electronics in ubiquitous computing. In *Proc. ACM CHI EA '12*, 2237–2242.
21. Pereira, T., Rusinkiewicz, S., and Matusik, W. Computational Light Routing. *ACM Trans. on Graph. 33*, 3 (2014), 1–13.
22. Pruijmboom, A., Schemmann, M., Hellmig, J., Schutte, J., Moench, H., and Pankert, J. VCSEL-based miniature laser-Doppler interferometer. In *Integrated Optoelectronic Devices 2008*, Intl. Society for Optics and Photonics (2008), 69080I–7.
23. Raskar, R., Beardsley, P., van Baar, J., Wang, Y., Dietz, P., Lee, J., Leigh, D., and Willwacher, T. RFIG Lamps: Interacting with a Self-Describing World via Photosensing Wireless Tags and Projectors. *ACM Trans. on Graph. 23*, 3 (2004), 406–415.
24. Rekimoto, J., Ullmer, B., and Oba, H. DataTiles: a modular platform for mixed physical and graphical interactions. In *Proc. ACM CHI '01*, 269–276.
25. Ren, P., Wang, J., Snyder, J., Tong, X., and Guo, B. Pocket reflectometry. *ACM Trans. on Graph. 30*, 4 (2011), 45.
26. Sato, T., Mamiya, H., Koike, H., and Fukuchi, K. Photoelastictouch: Transparent rubbery tangible interface using an lcd and photoelasticity. In *Proc. ACM UIST '09*, 43–50.
27. Sharan, L., Liu, C., Rosenholtz, R., and Adelson, E. Recognizing materials using perceptually inspired features. *Intl. Journal of Computer Vision 103*, 3 (2013), 348–371.
28. Shih, Y., Davis, A., Hasinoff, S. W., Durand, F., and Freeman, W. T. Laser speckle photography for surface tampering detection. In *Proc. CVPR '12*, 33–40.
29. Swiontkowski, M. F. Laser Doppler Flowmetry – Development and Clinical Application. *The Iowa Orthopaedic Journal 11* (1991), 119–126.
30. Taylor, B. T., and Bove, Jr., V. M. Graspables: Grasp-recognition as a user interface. In *Proc. ACM CHI '09*, 917–926.
31. Wiese, J., Saponas, T. S., and Brush, A. B. Phoneprioception: enabling mobile phones to infer where they are kept. In *Proc. ACM CHI '13*, 2157–2166.
32. Willis, K., Brockmeyer, E., Hudson, S., and Poupyrev, I. Printed optics: 3D printing of embedded optical elements for interactive devices. In *Proc. ACM UIST '12*, 589–598.
33. Willis, K. D. D., and Wilson, A. D. InfraStructs: fabricating information inside physical objects for imaging in the terahertz region. *ACM Trans. on Graph. 32*, 4 (2013), 138:1–10.
34. Xiao, R., Harrison, C., Willis, K. D., Poupyrev, I., and Hudson, S. E. Lumitrack: low cost, high precision, high speed tracking with projected m -sequences. In *Proc. ACM UIST '13*, 3–12.
35. Yang, X., Grossman, T., Wigdor, D., and Fitzmaurice, G. Magic finger: Always-available input through finger instrumentation. In *Proc. ACM UIST '12*, 147–156.
36. Zizka, J., Olwal, A., and Raskar, R. Specklesense: Fast, precise, low-cost and compact motion sensing using laser speckle. In *Proc. ACM UIST '11*, 489–498.