

Surface Material and Roughness Sensing Using mmWave via Surface Scattering and Ambient Vibrations

Yawen Liu*
Carnegie Mellon University
Pittsburgh, USA
yawenl@andrew.cmu.edu

Bert Shan*
Carnegie Mellon University
Pittsburgh, USA
bshan@andrew.cmu.edu

Swarun Kumar
Carnegie Mellon University
Pittsburgh, USA
swarun@cmu.edu

Abstract

In this paper, we explore a system to sense the roughness of surfaces, even if obstructed from the field of view. We pose this question in the context of robotic grasping and manipulation to explore whether robots can learn the texture of objects prior to grasping them. Importantly, we seek to do so in a completely contact-free fashion (ruling out tactile sensors), despite obstructions (ruling out cameras and lidar). We present mmTexora, a novel roughness sensing system using mmWave radar. mmTexora leverages ambient vibrations that produce temporal phase variations to objects in everyday environments, when perceived by radar. We demonstrate how these phase variations convey valuable information about the structure of bumps and ridges on a surface, thereby revealing details about surface roughness. We then develop a signal processing and deep learning pipeline that extracts surface roughness from the signal's temporal variations. mmTexora uses this information to classify surface textures when the material type of an object is known. Conversely, mmTexora can also classify material types for objects when they are known to have similar textures. We perform a qualitative study on a robotic arm tested on diverse objects, where either texture or material type is varied individually. Our classification model achieves an average surface classification accuracy of 93.7% on 50 surfaces that are commonly seen in daily life, with an average absolute error of 0.11 mm in roughness measurements.

CCS Concepts

• **Computer systems organization** → **Embedded and cyber-physical systems**; *External interfaces for robotics*; *Sensor networks*.

Keywords

Millimeter-wave sensing; robotic perception; haptic sensing.

ACM Reference Format:

Yawen Liu, Bert Shan, and Swarun Kumar. 2026. Surface Material and Roughness Sensing Using mmWave via Surface Scattering and Ambient Vibrations. In *ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys '26)*, May 11–14, 2026, Saint Malo, France. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3774906.3800459>

* Both authors contributed equally to this work.



This work is licensed under a Creative Commons Attribution 4.0 International License. *SenSys '26, Saint Malo, France*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2309-4/26/05
<https://doi.org/10.1145/3774906.3800459>

1 Introduction

In this paper, we ask: *Can we use mmWave radars to understand textured surfaces?* The two properties - material type and surface roughness - are important characteristics that directly influence a robot's ability to interact with an object. The material of an object determines key factors such as friction in grasping and manipulation. Even among objects of the same material, surface roughness provides valuable cues for interaction and manipulation. An object's roughness quantifies fine-scale differences in surface heights. For robotic grasping, roughness sensing enables robots to estimate the coefficient of friction and apply appropriate forces without causing damage. For human-tool interaction, roughness provides intuitive tactile feedback for operation. Applications range from cognitive research [7, 8], immersive gaming [27, 44], assistive devices [25, 38] to robotic manipulation [11, 36], and teleoperated operation [37]. A critical first step is to identify surface properties, including roughness and material, to enable real-time feedback.

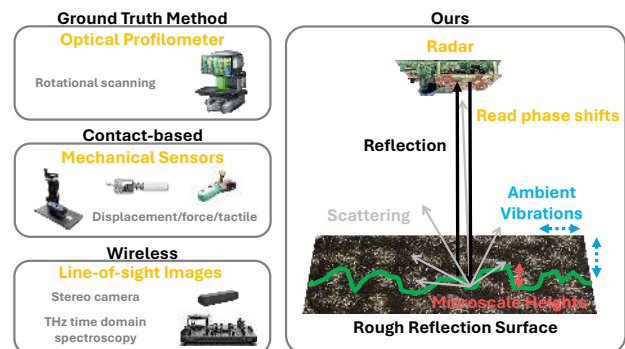


Figure 1: The local roughness can generate scatterings with phase shifts related to ambient vibrations and roughness.

While rich work exists on contact-based roughness sensing (e.g., torque [13], vibration [35], microphones [18]), these approaches inherently require physical interaction with the object. This adds latency and limits mobility, as such sensors cannot reach distant objects or hard-to-grasp objects. On the other hand, most non-contact methods, such as cameras and microscopes, fail in low-light conditions or when the line of sight is obstructed. In contrast, radio frequency-based sensing offers resistance to obstruction; however, prior work at sub-6 GHz (e.g., Wi-Fi bands) fails to resolve micrometer-scale roughness detection with limited resolution and is best limited to material classification [41, 42, 45, 48]. Recent work in sub-THz/THz spectrum offers promising roughness imaging capabilities [46, 49], yet they are bulky, expensive, and sensitive to




Figure 2: mmTexora's design overview.

environmental dynamics. This motivates us to explore how the mmWave spectrum (tens of GHz) situated between these frequency bands can perform accurate roughness sensing.

This paper proposes mmTexora, a surface sensing framework using commodity 77 GHz automotive mmWave radars. We demonstrate that, despite operating at a wavelength of a few millimeters, careful analysis of the radar data can capture micrometer-level waviness in surfaces, enabling accurate surface classification. Inspired by active vibrometry sensing, we develop a learning-based method that classifies one surface property: either material or roughness, when the other is given. Our approach leverages naturally occurring ambient vibration sources in the environment, such as building ventilation and structural vibrations, eliminating the need for an external actuator. We implement and evaluate mmTexora on 50 textured surfaces, demonstrating a classification accuracy of 93.7%, and an average absolute error of 0.11 mm relative to the optical baselines. To our best knowledge, we are the first paper to study two important surface properties, roughness or material sensing, through an integrated pipeline from a compact low-cost, single-modality platform (Table 1).

mmTexora's key insight stems from a curious phenomenon we observed: object reflections producing unexpected phase variations. This observation emerged from a relatively simple experiment – placing a single object on a static surface (e.g. table) in front of the mmWave radar and isolating the desired target. Given that our entire setup (radar, object, table, etc.) is static, we would anticipate that the phase remains constant over time. Surprisingly, however, we observed subtle temporal variations in the phase. We initially attributed this to system noise; we found that the phase variations were curiously systematic and correlated very strongly with the surface roughness of the object. However, this experiment was surprising in two respects: *why does the phase of a static object vary at all, and how does this variation relate to roughness?*

Through a series of subsequent experiments (described in Sec. 7), we conclude that the source of surface phase variations was due to the object itself vibrating owing to ambient disturbances – such as air vents, machinery, fans, etc. In hindsight, this observation is unsurprising. While the wavelength of 77 GHz mmWave radar is 3.9 millimeters, a displacement as small as 10 microns results in an observable phase shift of 1.85 degrees. A natural environment, such as a typical office space (even a quiet one), can easily produce such levels of displacement from ventilation alone. To validate this, we clamped the object carefully in an external vibration-free special environment (validated with a vibration sensor) and observed no phase shifts. This still leaves the second part of our question: *why*

do these vibrations and the phase shifts they induce correlate with the roughness of the object?

To explain the roughness-phase correlation, it helps to understand that the texture of an object ultimately depends on the height variations of the surface, as shown in Fig. 1. When a radar signal impinges on a surface, it scatters at different points on the object's surface. Each of these points lies at a slightly different distance depends on the precise location where it impinges. Object vibrations caused by ambient disturbances induce slight shifts in the mmWave signals' impinging point on the surface that vary subtly, and these small differences in propagation distance result in measurable phase variations. These small differences in propagation distance result in measurable phase variations. As an extreme example, a perfectly smooth surface vibrated parallel to the surface would not produce phase variations at all. In contrast, a highly jagged surface would induce phase variations that correlate strongly with the extent of height variations. Sec. 4 presents a series of validations for phase variations which are explicable with the roughness-related surface height variations measured with an optical microscope. Having established this correlation, we now ask our main question of interest – *how best can we process the radar signals to classify the roughness?*

Our key technical challenge is that the surface roughness influences the measured radar signal phase in complex ways under ambient vibrations. Given that these surfaces are vibrated in an uncontrolled fashion, the precise movement of the surface dictates the phase variations observed. In other words, the precise phase temporal variations from any given surface, even with a known roughness, are extremely hard to predict. Our key insight to resolve this is to study the *distribution* of phase variations over time rather than the precise temporal trend. We rely on a common approximation that the roughness of any textured surface follows a Gaussian distribution characterized by root mean square (RMS) heights along the vertical axis [10]. As a result, under diffuse scattering, phase shifts allow us to capture roughness properties.

The rest of the paper describes our technical approach to learn roughness features under ambient vibrations in typical environments with signal multipath. As illustrated in Fig. 2, we first perform background subtraction and utilize the minimum-variance distortionless response (MVDR) across radar chirps to suppress ambient multipath and strong vibrations of objects that are not of interest. We observe that most mechanical vibrations occur at low frequencies (0.5 – 5 Hz), allowing even short recordings of 1–2 seconds to capture sufficient surface responses. Sec. 5 and 6 detail the necessity of using a data-driven approach to surface classification. To capture subtle and complex temporal features that might be overlooked by traditional signal processing pipelines, we adopt InceptionTime [20], a deep learning architecture specifically designed for surface classification tasks.

Applications: mmTexora opens up several applications in the fields of robotics and human-computer interaction:

- *Surface Identification:* mmTexora enables accurate sensing of surface material and roughness, which can help automated quality control in smart manufacturing, allowing early detection of wear, defects, or contamination.

Reference	Modality	Coarse-grained Material Sensing	Fine-grained Texture Sensing	No Target Manipulation	Sensor Requirements	
					Contactless	Extra hardware Required
[23, 32, 47]	contacted-based	Yes	Yes	No	No	-
[21, 29]	active vibrometry	Yes	No	No	Yes	Speaker
[26, 40]	reflective tags	Yes	No	No	Yes	RF tags
[45, 48]	radio frequency	Yes	No	Yes	Yes	Drone/Motion stage
[46, 49]	THz spectroscopy	No	Yes	Yes	Yes	Beam focusing lenses
[41, 42]	mmWave Radar	Yes	No	Yes	Yes	-
mmTexora	mmWave Radar	Yes	Yes	Yes	Yes	-

Table 1: mmTexora exploits surface scattering to enable a single mmWave radar to sense tactile and material without manipulating targets or sensors. We highlight preferred features in green and undesirable components that hinder mobility in red.

- *Robotic manipulation:* mmTexora can help robots identify surface properties such as friction and texture, allowing robots to dynamically adjust grip force and motion strategy.
- *Augmented Interaction:* Surface-aware sensing can enhance haptic feedback, enabling richer human-computer interfaces.

Limitations: We emphasize a few important limitations of mmTexora (details in Sec. 8): (1) mmTexora is designed to classify one surface attribute at a time (either material or roughness) while the other remains fixed; we leave simultaneous sensing of both attributes through larger-scale training to future work. (2) mmTexora performance degrades under significant environmental changes; therefore, our approach explicitly emphasizes the incorporation of new calibration data into the training set to help the system adapt to ambient vibration patterns over days.

Contributions: The paper makes the following contributions:

- We introduce mmTexora, the first system that performs single-modality wireless sensing to identify either one of two key surface properties: roughness and materials.
- mmTexora pushes the radar object detection from simple material labeling to fine-grained roughness characterization.
- We build a data-driven model to perform surface identification and further predict surface roughness based on ground truth labels; mmTexora achieves an average classification accuracy of 93.7%, and an average absolute error of 0.11 mm in surface heights along 50 surfaces with mixed materials and roughness.
- We build a prototype system using commercial radars and evaluated its performance in a variety of scenarios, including a robotic grasping arm in a manufacturing facility.

We implement mmTexora on a commercial mmWave radar testbed (TI IWR1843Boost [1]) and evaluate system performance on 50 different textured surfaces. We perform microbenchmark studies to show the system’s resilience and robustness in Sec. 7. We demonstrate that mmTexora can learn new surfaces with a few calibration samples provided. We further train the network model to estimate surface roughness as height variations and compare with the optical profilometer’s ground truth measurements.

2 Related Work

We outline prior work on material and roughness sensing in Table 1.

Contact-based multimodal tactile sensing: Haptic sensing on textures and roughness plays a key role in human-computer interaction and robotic applications. Single modalities such as microphones

and friction sensors are used to infer roughness upon contact with a surface [18, 24, 28, 31, 33, 35]. Some work explores multi-modal processing, combining vision-based, microphones, and vibration sensors to improve classification accuracy [22, 23, 32, 47]. In this work, we explore an entirely contact-free wireless solution that avoids any damage or contamination of the targets.

Optical profilometer: Optical profilometer is a more advanced non-contact solution, which measures surface topography at sub-nanometer resolution [5, 17, 39]. The profilometer can illuminate a testing area with laser or white light to synthesize high-resolution 3D surface scans and quantify roughness as distance profiles along elevation and horizon. However, this equipment is specialized for laboratory use and is neither affordable (costs approximately \$5,000) nor portable for mobile or robotic systems. We use a Kayence VR-5000 profilometer [5] to collect ground truth data.

Radio Frequency based material sensing: Radio frequency (RF) systems for material detection fall into three broad categories: RFID tags, mmWave radar, and other commodity wireless technologies. RFID-based sensing [40, 43] measures frequency shifts from tags attached to the targets, whose impedance matching changes as the material’s dielectric properties differ. IntuWition [45, 48] uses commodity WiFi channel and polarization to identify material types. [41, 42] explores mmWave radar and range-azimuth readings to classify materials [21, 26, 29, 50].

RFVibe [29] is the closest to our work, which leverages *actively generated acoustic signals* to induce amplitude and phase changes in radar reflections from different materials. mmTexora differs from this work in two ways: (1) mmTexora uses ambient vibrations instead of active vibrometry with fewer constraints on setup; (2) mmTexora senses surface roughness while RFVibe senses material properties only. To the best of our knowledge, prior mmWave radar sensing systems across the board do not sense roughness. As summarized in Table 1, existing methods either largely overlook roughness features or depend on active motion, either from the sensors themselves or via manipulation of the objects being sensed. In contrast, our approach focuses on leveraging ambient environmental signals to infer surface characteristics.

Terahertz surface scattering: Terahertz spectroscopy uses broadband RF radiation (0.1-10 THz) with lenses to collimate and focus its short-wavelength (30 μ m to 3 mm) signals to sense surface height profile with line-of-sight captures [30, 49]. However, their

bistatic transceiver configuration is highly sensitive to incident angles, occlusion, and mechanical setup – even a displacement at the submillimeter level can introduce significant artifacts. In contrast, mmTexora aims to provide a comparable sensing resolution for mobile contexts with tolerance to interference and blockage.

3 Primer on Surface Sensing

3.1 mmWave Radar Sensing

mmTexora uses a mmWave frequency modulated continuous wave (FMCW) radar with 2 transmitters (TX) and 4 receivers (RX) to capture reflected signals from surfaces of interest. Specifically, we measure the signal amplitude and phase shift over a sampling period to characterize both the **material** and **texture**. The received signals at the n^{th} RX antenna and the chirp k^{th} can be denoted as:

$$Y_{n,k}(t) = \sum_{m=1}^{TX} \alpha * \exp \left(2\pi j f_o t + \pi j \frac{BW}{T} \left(t - \frac{T}{2} - kT \right)^2 \right) \quad (1)$$

, where α denotes the complex reflective coefficient and BW is the bandwidth of the chirps with starting frequency f_o and a duration of T at instant t . Conventional FMCW radar estimates target range using *Range FFT* amplitude; unfortunately, its range resolution is limited by bandwidth ($\frac{c}{2BW}$), up to 3.8 cm for our 77GHz radar. To achieve a finer resolution, we can exploit phase information at the k^{th} chirp: $\varphi = 2 * 2\pi \frac{d}{\lambda}$ traveling round-trip distance $2d$. The phase difference between the chirps reflects displacement from scattering or vibrations, with a resolution limited by the chirp rate: $\Delta\varphi(k, k-1) = \frac{\Delta\varphi(k, k-1)\lambda}{4\pi}$. Consequently, relying solely on phase information may misrepresent vibrations that occur faster than the sampling rate or evolve more slowly across multiple data frames.

3.2 Radar Scattering Fundamentals

Here, we discuss some basic concepts on how mmWave signals scatter from textured surfaces. We ask the following question. How does surface roughness influence its reflectivity?

Surface properties and scattering: While material sensing has been well studied [41, 42, 45, 48], mmTexora also investigates textures. We consider roughness as minute variations in surface height on the microscopic scale, as shown in Fig. 3. Mathematically, roughness is characterized by the standard deviation of horizontal variation (RMS height) from the lowest horizon on the surface. Typical machinery tools require smoothness with height irregularities less than 5 μm . Fabrics and furniture usually have h_{rms} ranging from ten to a hundred micrometers. Rough surfaces that human fingers can feel or visually identify as rough, such as stones and carpets, have h_{rms} in hundreds to thousands of micrometers [3]. We hypothesize that these microscopic height variations induce diffuse scattering, as incident waves are dispersed in multiple directions. We will analyze and evaluate this effect in detail in Sec. 4.

Electromagnetic Roughness: According to the Raleigh scattering criterion [10] a surface is considered *electromagnetically rough* if its RMS height exceeds approximately $\frac{\lambda}{8}$ (e.g. occurs if $h_{rms} < 0.49\text{mm}$ for 77 GHz radar), causing significant diffuse scattering of incident waves. Fig. 3 shows the scanned image of a rough carpet obtained with an optical profilometer, where the heat map represents the height profile (Fig. 4) – a typical electromagnetically rough surface.

4 Understanding Surface Roughness and Radar

In this section, we investigate how rough surfaces affect mm-Wave radar responses, considering that objects undergo ambient vibrations. Previous research [15, 16, 30, 49] has investigated electromagnetic wave scattering and roughness through the integral equation model (IEM), where the surface's RMS heights exhibit a Gaussian distribution (Fig. 4). Specifically, we leverage the IEM to estimate the reflected waves as a function of *signal frequency*, *incident/receiving angles*, and *RMS height distribution*. We simplify this scattering model to better fit monostatic radar (i.e., co-located transmit and receive antennas), where we assume the receiving angles equal the incident angles: $\theta_r = \theta_i$ and $\phi_r = \phi_i$ for elevation and azimuth, respectively (Fig. 3). This simplification facilitates the interpretation of surface scattering effects from roughness, which we summarize as being primarily dominated by ambient vibrations and characterized in two dimensions: signal phase and amplitude.

4.1 Signal Phase and Surface Roughness

We first explore *ambient vibrations* on the phase of mmWave surface scattering using proof-of-concept experiments and explore its correlations with surface texture. We present a series of experiments to reveal the following key insights.

Ambient vibrations subtly alter radar phase: Ambient disturbances cause objects to vibrate slightly, leading to subtle shifts in the point where mmWave signals strike the surface. Environmental vibrations typically exhibit low-frequency periodicity of 0.2-2 seconds (10 radar frames for our configuration). A full micro motion cycle spans tens of chirps (about 64 μs per chirp and 20 ms per frame with idle periods between chirps). This implies that we need to observe multiple chirps, or frames that contain a pack of chirps, to visualize the influence of ambient vibrations on radar.

Temporal phase data captures information on texture: Recall from Sec. 3.1 that the phase change is proportional to the level of micro-motion displacement induced by ambient vibration and scattering. *In an ideal world without vibrations, we should expect an absolute constant phase value versus time in a perfectly static setting, even from a surface that can induce strong diffuse scattering.* In practice, however, subtle environmental vibrations cause dynamic scattering at different points on the surface, leading to time-varying radar phase responses. We formulate the phase changes from ambient vibration and diffuse scattering as follows:

$$\Delta\varphi(t) = \frac{4\pi\Delta h_{rms}}{\lambda} A_{vib} \sin(2\pi f_{vib}t) \quad (2)$$

, where Δh_{rms} is the height change on the microscale due to the variation of the scattering point on the surface, and $A_{vib} \sin(2\pi f_{vib}t)$ approximates the vibration of the ambient environment at a certain strength A_{vib} and time t . Although it is clear that signals do indeed capture a texture-dependent component (i.e. h_{rms}), it remains a challenge to isolate it from all other extraneous factors (e.g., environment/vibration-dependent terms) – the main problem we tackle in this section.

Multipath phase components combine within chirps: While the IEM test can predict the power of the scattered signal, it does not capture phase information. To remedy this, we assume that scattered paths interfere *consistently* within a chirp and are *uniquely*




Figure 3: Scattering on a rough carpet.




Figure 4: Heights distribution from scans.




Figure 5: Exploratory studies testing samples: comparing amplitude and phase changes with different clutter, materials, and roughness.

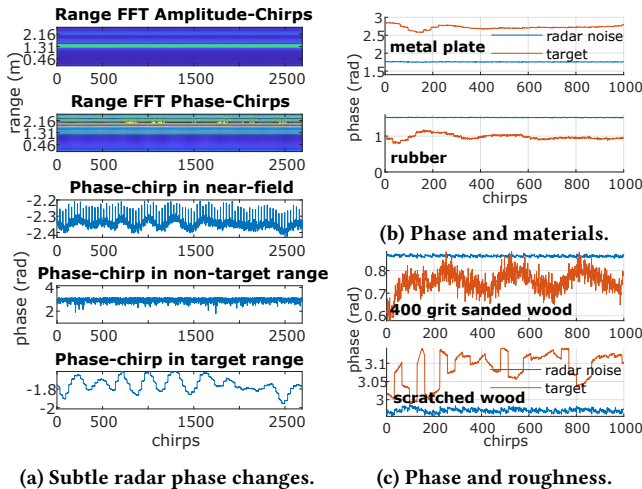


Figure 6: Radar unwrapped phase response: left shows the target's phase change, right shows the impact of surface properties and ambient vibrations.

influenced by the height distributions. Such diffuse scattering from surface texture causes each point to contribute a distinct phase delay. Since the chirp duration (64 μ s) far exceeds the total propagation delay of multipath signals (a signal can travel 9000 meters in that time), the received phase effectively sums up these delays. This effectively creates a direct link between the signal's phase components and the surface height distribution – i.e. roughness.

We visualize this observation in Fig. 5, which shows that both external vibrations and surface scattering contribute. In fact, they *modulate* the radar echo simultaneously, capturing both surface texture and ambient clutter. Fortunately, we find *although clutter vibrations dominate absolute phase and amplitude shifts, the relative phase and amplitude variations preserve texture-related features*. For example, temporal phase variations are more pronounced and stable for a rough wooden surface than for a smooth surface (Fig. 5 and Fig. 6c). We build on this insight to develop our texture feature identification in Sec. 5.

Ambient vibration induced phase shifts are observable well-above phase noise: A natural question is whether the phase shifts induced by ambient vibrations are detectable above the noise, such as thermal noise or hardware imperfections. Could these noise sources alone account for the unexpected phase variations, overshadowing the effects of surface scattering? In fact, Fig. 6a validates that our radar's phase noise is far below the phase introduced by

the ambient vibrations. We observe the near-field phase response, demonstrating high noise from spherical propagation, while a far-field range bin with no objects exhibits minimal phase changes. Meanwhile, phase measurements are largely stable in locations further away, albeit with negligible levels of thermal noise. We note that careful radar configuration, tuned to the periodicity of ambient disturbances, is crucial for accurately capturing phase variations. Specifically, mmTexora uses frames of 32 chirps, each 64.13 μ s long with 10 μ s idle between chirps and a frame duration of 20 ms.

Roughness-to-Phase correlation remains consistent: In diffuse scattering, reflections from small surface wavinesses cause phase changes even within a single range bin (Eq. 2). By comparing the unwrapped phase between materials, we find that the phase variations remain consistent in the same clutter, as shown in Fig. 6b. Interestingly, we observe that smooth surfaces are slightly more vulnerable to changes in clutter. Since the phase is sensitive to sub-wavelength displacements, scattering tends to augment the randomness and decorrelation in the unwrapped phase pattern in Fig. 6c. This finding implies that the fine-grained texture changes in the target are traceable.

4.2 Signal Amplitude and Surface Roughness

In this section, we show that beyond the signal phase (our focus thus far), signal amplitude also proves pivotal in learning the surface.

Surface material is highly correlated with RF reflectivity, despite vibrations: As we described in our primer (Sec. 3.2) and prior material-sensing papers [29, 45, 48], RF reflectivity measured from amplitude is strongly influenced by the material type. Interestingly, this amplitude vs. reflectivity relationship is preserved, even in the presence of both subtle and strong variation in the environment (Fig. 5). In Fig. 7e, we observe that the higher-order multipath bounces appear as weaker peaks in further range bins, and in Fig. 7f, we observe that rough surfaces produce stronger higher-order bounces than smooth ones, as they tend to diffuse power more sideways. This is consistent with prior results [29, 45, 48] since ambient vibrations are unavoidable. It is unsurprising, given the fact that ambient vibrations influence signal phase much more than amplitude. We next show that we can, in fact take advantage of these subtle amplitude *changes* to detect texture (rather than material) as an addition to signal phase.

Roughness impacts amplitude changes, not its absolute value: Conventional radar-based material sensing focuses on analyzing absolute Range FFT amplitude to infer reflectivity. In contrast, we

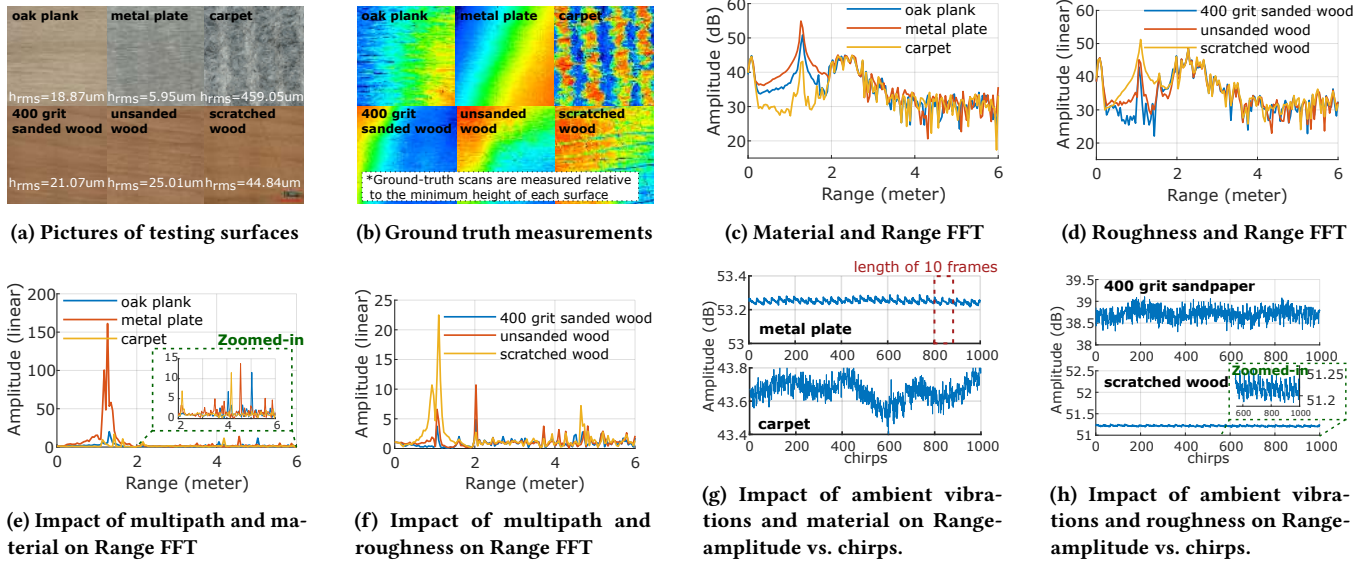


Figure 7: Impacts of surface properties on signal power.




Figure 8: The scattering coefficient analysis reveals the reflected power across the space under the radar's view.

aim to find the relative temporal-amplitude variations that stem from radar scattering changes. Fig. 7c and 7d show the Range FFT comparing the textures of materials with different roughness. Our observations show that both the material type and the surface texture significantly affect the reflectivity. Notably, when comparing wooden samples that only vary in roughness, we find that diffuse scattering leads to rapid changes in reflected amplitude. This is because rough surfaces tend to occasionally scatter back more power when paths constructively interfere as the object vibrates. This behavior explains why previous work [29, 45, 48] could only support coarse material classification but struggle with textures: *textures add ambiguity to the Range FFT, hindering the classification task that simply utilizes reflected power from a single time snapshot.*

Roughness-to-amplitude correlation can be enhanced with carefully chosen radar parameters: Next, we discuss how various parameters of the radar system matter. We observe that mmTexora is minimally affected by signal frequency variations from FMCW chirps. To evaluate this, we performed the IEM test [15, 16, 30, 49] using a known $h_{rms} = 1mm$ and estimated the scattering behavior. Fig. 8a illustrates the radar's beam pattern, and Fig. 8b simulates the normalized signal power received from a point reflector along the elevation and azimuth axes. We then incorporate the surface height parameters (i.e. $h_{rms} = 0.5mm$), where Fig. 8c shows the

normalized reflected power of a textured surface along the two axes. We observe that the varying frequency of chirps has a negligible impact, except at extreme incident angles. In Fig. 8d, we map the scattering profile to the radar's reflectivity test in Fig. 8b, which suggests that placing the test sample near the center of the radar field of view yields the best performance. As highlighted in Fig. 8e, where we compare the measured results of a small surface of $h_{rms} = 0.61mm$, it has a matched power distribution with simulations. Thus, to achieve good detection, we first apply conventional range-azimuth methods to locate the object and maximize signal quality.

5 Signal Processing for Feature Robustness

Resilience to noise, interference, and clutter limitations: We find that projecting target data onto a reference background can help isolate the target and reduce multipath effects when captures are close in time. Specifically, we record two measurements, one with the surface and one without, for cancellation, but these dynamics do not *always* cancel out. We empirically find that reference capture is helpful only if taken within a short period, approximately 10 minutes; otherwise, the ambient vibrations change fundamentally and would further pollute the data. We discuss the impacts of these factors in Sec. 7.6. In addition, we leverage the radar channel

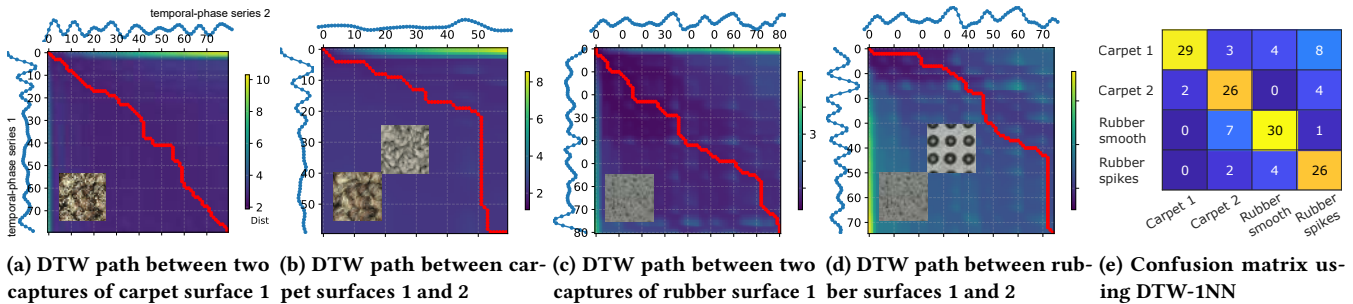


Figure 9: Exploratory studies to identify time series phase changes are unique features of surface roughness.

to further zoom in on the target surface. Our radar provides 8 virtual channels (2 TX and 4 RX). We apply the Minimum Variance Distortionless Response (MVDR) algorithm to optimize the beam, focusing on the surface while suppressing distortion and preserving the target phase with minimal interference. We then isolate the temporal phase response by subtracting the MVDR-enhanced amplitude from the original range FFT output, increasing the sensitivity to subtle variations within the target’s range bin.

Temporal feature resilience to vibrating clutter: Building on observations in Sec. 4, we postulate that the radar Range FFT’s temporal changes are directly related to texture. However, the most critical challenge mmTexora faces is to effectively extract surface scattering information when it is buried far below the contributions of ambient vibrations. Inspired by speech recognition, which identifies words despite variations in voice and pitch, we aim to extract features from surface temporal data for classification tasks, treating ambient vibrations as analogous to noise in tone and pitch. From there, we attempt Dynamic Time Warping (DTW) as a proof-of-concept metric to guide the design of our classification algorithm in Sec. 6. DTW paths can measure the similarity between two temporal sequences that can be misaligned [6, 12, 14]. Interestingly, we observe a stronger alignment, with paths lying closer to the diagonal, for two radar captures of the same surface in Fig. 9a and Fig. 9c. In contrast, even visually similar carpets in Fig. 9a show distinct DTW paths. To demonstrate further, we run a small-scale nearest neighbors (DTW-1NN) classification on two carpets and two rubber surfaces with data from different days and locations (Fig. 9e). This result confirms that the temporal data retain meaningful roughness and material cues, despite unpredictable ambient vibrations, making this problem well-suited for a data-driven solution.

6 Classification and Roughness Model

6.1 Time-series Deep learning model

Why we need a data-driven approach? In Sec. 5, we have shown that traditional time series similarities such as DTW can help identify the hidden clues of the temporal-amplitude and temporal-phase features related to surface properties. *However, is the problem solved with a simple DTW-1NN classification? The short answer is no.*

Choosing the optimal model: Unfortunately, the simple DTW-1NN model struggles to scale due to its inability to capture hidden features behind clutter noise, its high computational cost, and its




Figure 10: InceptionTime network architecture for mmTexora time series classification with simplified bottleneck layer.

susceptibility to learning hidden patterns in the presence of varying ambient vibrations, making it unsuitable for surface roughness prediction. We can observe notable detection errors in Fig. 9e. Therefore, we implement a data-driven approach that can better extract hidden features with larger datasets. Although time-series data are inherently challenging for extracting physics-based insights in vibrometry, we prefer InceptionTime over recurrent neural networks and long short-term memory models due to its ability to simultaneously capture both short- and long-term features [20]. One unique advantage is that InceptionTime does not require careful tuning and saves computational resources with parallelized convolution layers, which instead is filtered and parameterized based on the data sampling rate. This feature adds robustness against overfitting and vanishing gradients; thus, we choose InceptionTime.

6.2 Neural Network Architecture

Owing to the inherent challenge of modeling physics-based discriminative features from the vibrating clutter, we choose InceptionTime architecture to extract the hidden features primarily through a data-driven method. InceptionTime model is a convolution-based set of multiple Inception modules that learns features simultaneously [20]. From Sec. 3, we have concluded that the temporal trends in amplitude and phase can reflect surface properties, however, ambient vibrations might obscure these subtle features. To address this problem, we leverage InceptionTime filters [19], transforming the classification model’s inputs to a concatenation of multivariate time series with reduced complexity.

Fig. 10 illustrates the architecture of the model. Specifically, we first normalize the amplitude and phase time series obtained from the radar digital signal processing pipeline. We applied a sliding-window approach with smoothing and local peak detection to highlight temporal trends. These filters are independently and automatically parameterized to the temporal data based on radar



Figure 11: Experiment setup and environments




Figure 12: Experiment surfaces; Index of surface stays the same for all evaluation figures

chirp configuration, transforming them into a concatenation of a 4-channel final input with normalized phase, normalized amplitude, and filtered local peaks and troughs.

To further mitigate the perturbations from unknown vibrations, we introduce a bottleneck layer with a 1D convolution to reduce the dimensionality of the multivariate time sequences. Three convolution kernels with 10, 20, and 40 time steps are used to extract features in parallel. This allows the network to effectively capture both short-term and long-term features. To further enhance the stability, a stride one MaxPooling layer with a pool size of three is applied, followed by a 1D convolution. The output of parallel convolutions and MaxPooling is concatenated at each module. The cascaded Inception modules can extract hierarchical features, integrating with conventional layers for more effective residual connections. The trained results finish with an average 1D pooling, and a softmax dense layer for predictions.

7 Evaluation

7.1 Evaluation Setup

Hardware and data collection: mmTexora uses the 77 GHz TI IWR1843BOOST front-end (2 TX and 4 RX antennas used) with the DCA1000EVM [1, 2] evaluation board. Data were collected through mmWave Studio and Python socket connections.

Experiment setup: We evaluated mmTexora at 5 locations across 2 different buildings, including a typical office setup and a use case study in a manufacturing facility with a Yaskawa GP4 industrial robots [4] for grasping, as shown in Fig. 11 and Fig. 20a. We conducted data collection over several weeks to ensure diversity

from clutter and ambient changes. The target surfaces were placed between 0.2 m to 2 m from the radar in four setups: a tripod, a presentation board, a table, and a robotic arm manufacturing platform. Each data capture consisted of 80 frames in 1.2 seconds.

Surface properties: We collected a large data set that includes 50 different surfaces, spanning 11 material types and varying roughness (Fig. 12). All surfaces were standardized to 8×11 inches for direct comparison. To isolate roughness from material, we crafted multiple wooden squares from a wood panel (surface 36 in Fig. 12) with different levels of roughness using sandpaper and tools. This allows us to evaluate whether mmTexora identifies the true texture of the surface rather than the simple composition of the material.

Evaluation standards: We design mmTexora to perform classification on a single property at a time, where the output corresponds either to a surface category label or to an estimated surface roughness value. However, the sensing precision and model capacity are sufficient to jointly infer multiple surface properties simultaneously. To validate this capability, we performed experiments on a data set that includes (i) surfaces with different materials but similar texture or roughness characteristics and (ii) surfaces composed of the same material but showing different levels of roughness in Sec. 7.7. The classification performance remained consistent across both conditions, demonstrating the system’s ability to measure one of material types or roughness, when the other remains fixed.

Dataset and metrics: To facilitate training, we obtain a total of 5,270 radar captures. Our main evaluation uses 70 percent of the data split for model training and the remaining for testing, in a non-overlapping manner. We trim all data captures to the same frame length so that our InceptionTime model has a consistent number of dimensions for all input time series. We first present the overall accuracy (in %) and classification matrix (CM) results for detecting 50 surfaces, types of materials, and levels of roughness in h_{rms} measurements.

Ground truth and baselines: We use ground truth measurements from a Keyence VR-5000 optical profilometer [5] to label the roughness metrics, h_{rms} , in micrometer scale. In the rest of the evaluation section, we train and test the network 20 times for classification tasks and 10 times for numerical roughness estimates with different training and testing sets to verify stability and robustness. The results are compared with the baseline work [29] originally developed for material detection¹ and ground truth measurements from the profilometer. We did not use a speaker like that in RFVibe [29], but only their feature extraction and classification pipeline on our data.

7.2 Surface Classification

Experiment. We first present the overall performance of mmTexora in three classification tasks: (1) identification of 50 textures from the training dataset; (2) detection of material types with 11 classes, metal/non-metal/RF absorbing labels, and binary hard/soft labels; (3) coarse-grained roughness level labeling. Training and testing data were randomly selected from a mix of experiments at different times and locations. For material detection, we group the

¹We add the caveat that while these systems were not designed for texture, we simply explore the performance to sense texture with their material sensing pipeline – because we lack prior art performing radar-based texture sensing as baselines.

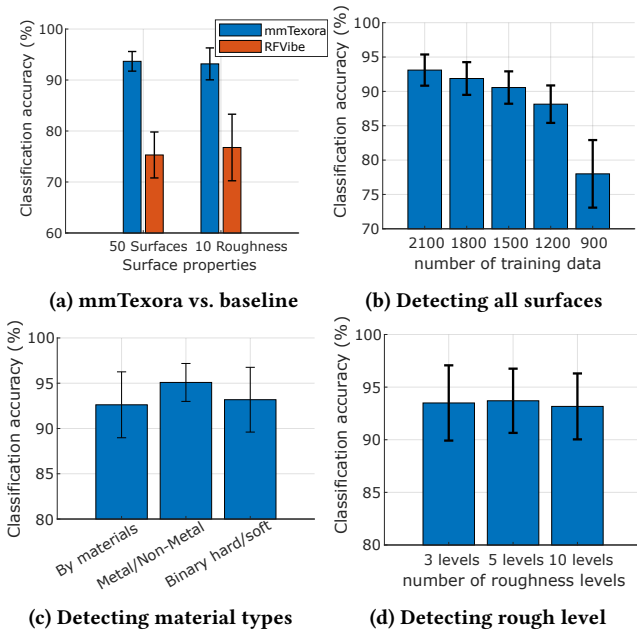


Figure 13: Performance on different detection tasks.

50 textures in 11 classes, including paper, plastic, and plastic fabrics (coated with reflective, shining surfaces), regular fabrics, leather, rubber, wood, foam, flooring materials, metal, and knitted surfaces (metal grids, silk, and tablecloth). When labeling texture features with roughness levels, we divide them according to the distribution of h_{rms} into 3, 5, and 10 sub-levels.

Result. mmTexora achieves an average accuracy of 93.7% classifying 50 surfaces, which is 18% higher compared to the baseline as shown in Fig. 13a. Fig. 14 shows an example of confusion matrix results sorted by material types. Since RFVibe [29] leveraged a similar phase feature through an active, controlled vibrometry method using a speaker, they provide some capability to generalize to ambient vibrations and can distinguish some surfaces with an average accuracy of 75.3% classifying 50 surfaces and 76.8% estimating roughness labels, but still significantly lower than mmTexora’s performance. Since RFVibe leverages the same temporal-phase features, it can distinguish between some surfaces of the same materials and different textures from the ambient vibrations. *In contrast, mmTexora’s main advantages are that it does not require active vibrometry like RFVibe and can eliminate the need for additional hardware (e.g., a speaker), while still maintaining robust performance under ambient clutter and vibrating sources.*

We observe that mmTexora outperforms the baseline in both material and roughness detection. As shown in Fig. 13b, the performance still achieves an accuracy of 78.0% with 6 training captures for each surface. mmTexora’s performance in detecting material types undergoes multipath and environmental variations, can still provide a promising accuracy of 92.6%, 95.1%, and 93.2% for three different material characterization tasks in Fig. 13c. *This results lead to another advantage of our design: mmTexora with a simple InceptionTime network can generalize to new environments and unseen surfaces with a few calibration data. In comparison, RFVibe, which uses a convolutional model that requires careful tuning of weights.*

7.3 Material and roughness classification

Experiment. We evaluate mmTexora’s classification performance among the two surface properties, showing confusion matrix results from the overall performance in Sec. 7.2.

Results. Fig. 15 illustrates the confusion matrix among all 50 surfaces; Fig. 16 illustrates the confusion matrix of all surfaces into 10 different levels of roughness. We observe in general, promising accuracy of more than 90% for detecting materials. When we separate the surface roughness into levels, we observe that only one level, between 227 and 337 μm yields our poorest (worst-case) accuracy of 84%. This is caused by the abundance of surfaces that fall into this range compared to other levels. In fact, the more different the surface roughness, the higher the accuracy, because the radar responses can be easily distinguished.

7.4 Roughness detection on surfaces made of the same material

Experiment. mmTexora’s main goal is fine-grained surface roughness estimation. To ensure we are not characterizing the minute difference in the surface material, but are truly extracting surface textures, we evaluate mmTexora’s performance in two different ways: categorizing the surface roughness into labels for classification, and predicting the numerical h_{rms} measurements as described in Sec. 4. We chose one material and created different surface roughness with the same size and weight as mentioned in Sec. 7.1 and followed the same experiment setup in Sec. 7.2.

Results. Fig. 17a shows the confusion matrix on detecting surface roughness on wooden planks. We can observe that the sandpaper and tool-scratches successfully created significant variances from their ground truth scans in Fig. 20b. We observe that most samples can be identified with high accuracy, with the sanded pieces with a 75% accuracy. This challenge may arise because the hand sanding was uneven, resulting in irregular surface roughness. Similarly, we also test mmTexora on rubber with different surface textures as labeled in Fig. 12. The confusion matrix in Fig. 17b show both rubber 30 and rubber 32 have excellent detection accuracy with their unique patterns. Although rubber 34 and rubber 48 are visually dislike in their textures, rubber 34 only achieves 67% accuracy because it is a thin layer of soft rubber, which can easily bend.

7.5 Numerical predictions on roughness

Experiment. To predict surface roughness numerically, we extend the network model in Sec. 6 by adding a 16-unit layer using a rectified linear unit (ReLU) activation function followed by a single-unit linear output layer. We calculate the loss as the mean squared error between the measured RMS height and the output estimate of the roughness measurements. We evaluate the model with 30 testing samples on each surface and compared their estimated mean and standard deviation with ground truth measurements.

Result. Fig. 18a shows the RMS height estimates for all testing data split from the 50-surface dataset. We observe that the estimations have a good match with the measured ground truth; nevertheless, the error increases for surfaces with smaller RMS height. Unfortunately, this is an inherent challenge because of the limitations of

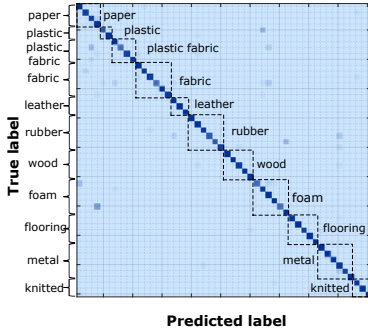


Figure 14: CM by all surfaces

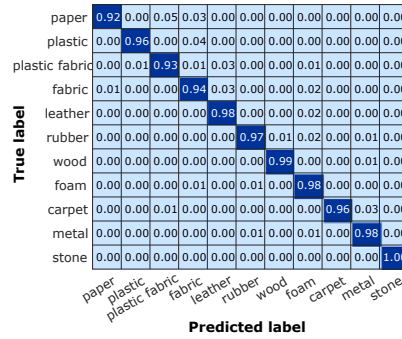


Figure 15: CM by material categories

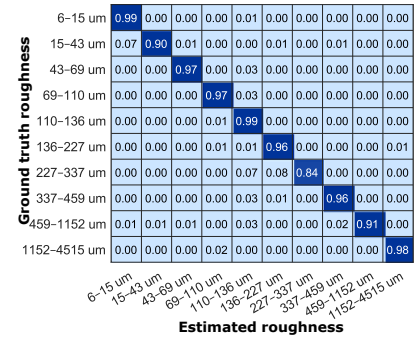





Figure 16: CM by level of roughness




(a) Confusion matrix of wood with different roughness



(b) Confusion matrix of rubber with different textures



(c) Compare the wooden surfaces from the same piece of plank

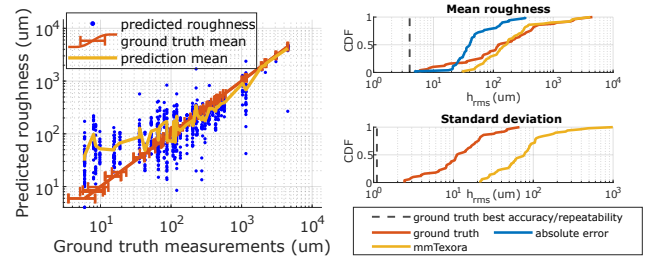


(d) Compare rubber surfaces with different textures

Figure 17: (a) - (b) Confusion matrix of roughness detection of surfaces made of the same material. (c) - (d) Predicted h_{rms} as roughness measurements compared to ground truth.

manual experiments - the error of ground truth measurements also increases for smooth surfaces with tiny $h_{rms} < 10 \mu m$.

Compared to previous work using THz spectroscopy [46], which reaches an absolute error of 0.039 mm in RMS height estimation, mmTexora achieves a median absolute error of 0.041 mm in RMS height estimation and an average absolute error of 0.069 mm. We find that mmTexora achieves comparable performance, while it can significantly reduce the costs of the hardware sensor. We note that mmTexora achieves good estimation accuracy for the surface with $h_{rms} > 0.1mm$, as shown in the CDF results of all 50 surfaces in Fig. 18b. Due to the wavelength limitation of the mmWave signals, our system cannot achieve the same level of repeatability as optical microscopes (ground truth methods have 4.0 μm accuracy and 1.0 μm repeatability [5]). mmTexora's estimations on the 50 surfaces dataset yield a median of 77 μm standard deviation. We believe this is still a valuable performance, considering that our mmWave radar has a wavelength of 3.9 mm. Hence, mmTexora can detect surface



(a) Roughness predicted h_{rms} vs. (b) CDF of predicted and measured roughness of all surfaces

Figure 18: Numerical estimations on roughness h_{rms} of all 50 surfaces and the comparison to ground truth.

features 50 times smaller than the radar's wavelength. We calculate an average absolute error of 0.11 mm and a median absolute error of 0.0398 mm for the estimations on the 50 surfaces, with the smallest absolute error of 7 microns from ground truth.

7.6 Microbenchmarks

(1) **Impact of surface area.** It is important to see how surface area sizes can alter system performance. In Fig. 19a, we compare three sizes, with the four surfaces with samples cut into squares of 15 cm, 20 cm, and 25 cm, and fit to our model trained in the main results. We observe that 8-by-11 inch (20.3-by-27.9 cm) samples achieve the highest precision, since they match the surface size of the training data. The overall test precision fluctuates less than 4.9%. A larger area leads to better detection performance, while a smaller area leads to greater variability and a larger standard deviation.

(2) **Impact of calibration data on roughness estimations.** We also perform roughness estimation experiments on testing samples of different sizes in Fig. 19b. Four surfaces with divergent roughness are used as testing data to find the absolute error in h_{rms} before and after calibration. The solid lines show the mean absolute error on the model tested without calibration data; the dashed lines show the results after calibration on 6 data samples. We observe smaller absolute errors after calibration; using a larger surface area sample also reduces this error. We calculate a mean absolute error of 0.174 mm and 0.076 mm before and after calibration, respectively.

(3) **Calibration to occlusion and changing clutter.** Owing to the reliance on feature extraction and interference with ambient clutter, mmTexora requires calibration captures to generalize to unknown

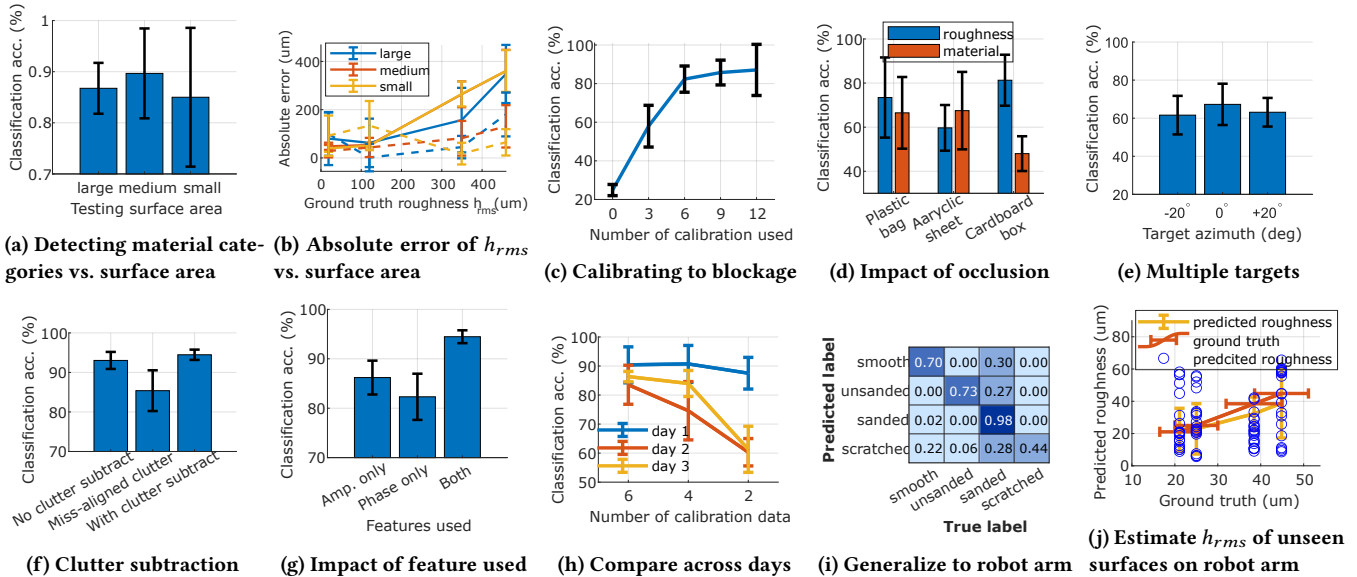


Figure 19: Micro benchmark studies on mmTexora under different parameters and scenarios: (a) impact of testing surface area size; (b) roughness estimations vs. surface area with and without calibration data; (c) impacts of calibration captures on obstructed objects; (d) impacts of different occlusions; (e) detecting multiple surfaces; (f) impacts of background cancellation; (g) compare individual mmTexora features; (h) compare performance with data collected across days; (i) generalization to unknown environment and known surfaces; (j) generalization to unknown environment and unseen surfaces.

surface sizes or clutter changes. On top of the trained model on the same surfaces but fixed sizes and no occlusions, we vary the calibration data size to find the minimal number of captures required to achieve effective performance. We place a total of 8 surfaces (used in Sec. 7.4) behind the occlusion objects as shown in Fig. 11. As in Fig. 19c, we observe that introducing 3 calibration captures reaches 58.0% accuracy, and using more than 6 calibration captures yields 82.3% accuracy, comparable to tests without occlusions. Similarly, we also compare calibration results for different days in Fig. 19h and conclude that the preferred number of calibration samples is 6.

(4) Impact of occlusion types. For vision-based and contact-based sensing methods, loss of line-of-sight data collection can lead to complete system failure. To verify performance change in non-line-of-sight conditions, we place our objects of interest behind a plastic bag, an acrylic sheet, and a cardboard box (Fig. 11), and then classify the material and roughness levels. We calculate an average accuracy of 73.4%, 59.7%, and 81.3% to identify roughness levels (up to 10 levels) and 66.5%, 67.5%, and 48% to identify materials behind the three occlusions, respectively, in Fig. 19d. This result matches our expectation, as obstructions attenuate the signal power and introduce additional multipath. Interestingly, we found that roughness detection is more robust than material detection. We believe that this is caused by the fact that the temporal-phase feature is closely correlated with the surface scattering as described in Sec. 4 – unwrapped phase does not change substantially when signal power drops, making it less influenced by the obstructions.

(5) Impact of multiple objects and placement. To test mmTexora with multiple objects presented in the radar field-of-view and to use beamforming techniques to isolate each object. We evaluate classification accuracy *directly without any calibration steps* with features extracted from the corresponding Range-FFT bins of each

surface using the trained model with 50 surfaces. We consider a total of 6 combinations, placing three surfaces at different azimuth angles as in Fig. 11. Fig. 19e shows an average classification accuracy of 61.6%, 67.3%, and 63.1% for objects at -20° , 0° and $+20^\circ$ from the radar’s perspectives. We show mmTexora has some resilience to multiple targets and find that the objects in the center exhibit the highest detection performance, as expected from Sec. 4.2.

(6) Impact of individual features and background cancellation. We discuss the importance of background cancellation and two individual features, amplitude and phase, in Sec. 5. We evaluate mmTexora with and without background cancellation, and then repeat the same experiment on individual features. Fig. 19f shows that subtracting a misaligned background can pollute the data and degrade the performance. When conditions permit, clutter subtraction leads to a small increase in accuracy of 1.4%. This implies that clutter subtraction is not always required if the model has been trained well in known environments. According to Fig. 19g, we observe that using both features outperforms individual features by 8.3% and 12.2% for amplitude-only and phase-only classification.

(7) Generalize to unseen environments and surfaces. Fig. 19i shows a confusion matrix when we deploy mmTexora to a robot arm, but train the model with data measured elsewhere. We find that mmTexora can generalize to unknown environments and unseen surfaces, albeit with some performance loss. Fig. 19j predicts the 4 rough wood planks (surface 36 is excluded from training data) using a model trained with the other wooden surfaces in Fig. 12 35, 37, 46, and 38. Although the absolute error increases to 0.0291 mm, mmTexora can still predict the roughness at the correct magnitude and scale. This demonstrates mmTexora’s potential to predict surface roughness even when the exact surface has not been trained, allowing similar surfaces to be used for generalization.

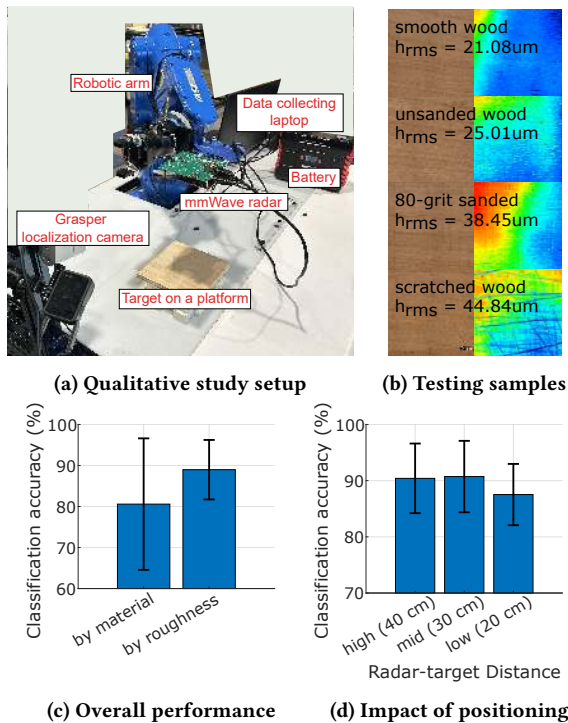


Figure 20: Qualitative study deploying mmTexora in an industrial environment and commercial robot arm grasping objects using a pre-trained model with three calibration attempts.

7.7 Qualitative Studies on Robotic Arm

Real-world use case study. To access the real-world performance of mmTexora, we perform experiments in a factory-like setting on a grasping robotic arm, where we find that such a manufacturing facility contains richer ambient vibration sources. Subsequently, we leverage 3 calibration captures to adapt the model trained from the quieter office setup. We test on surfaces 13, 14, 32, and 36; in addition, we test the four wooden planks with different roughness (Fig. 20b). We mount the radar to a robot arm, facing the operating platform, as shown in Fig. 20a.

Result. The results in Fig. 20c indicate the robustness of mmTexora in a new environment with 80.6% and 89.0% accuracy in detecting materials and roughness, separately. We further evaluate detecting all eight surfaces: four materials and four wooden pieces with different roughness levels, to evaluate performance with *simultaneous material and roughness sensing* based on the robotic arm’s positioning. As the robotic arm’s joints constrain its range of motion, we then evaluate performance at three radar-target distances (Fig. 20d). mmTexora achieves an average accuracy of 89.6% with only minimal variation across distances. Indeed, material and roughness characteristics are important for these robot-tool interactions, helping them grasp the item with the correct amount of force and position. Our results show that, although strong clutter vibrations can reduce performance, the system remains effective with environmental learning.

8 Discussion and Future Work

Practicality and Limitations: We acknowledge that the practical implementation of the proposed system faces several key challenges: (1) Our evaluation in Sec 7 considers two different scenarios for potential applications, including office and factory. However, our approach assumes relatively stable environmental conditions. An important assumption is that vibrating noise in daily life is often predictable [9, 34]. mmTexora requires a recalibration if the environment undergoes significant structural changes due to its sensitivity to clutter noise. Despite these limitations, we demonstrate robustness over time with continuous training in Sec. 7.6, suggesting that adaptive learning techniques can improve system robustness to evolving environmental conditions. (2) The system is currently optimized for objects within primarily indoor sensing scenarios; for instance, we evaluated mmTexora’s performance on a fixed robotic operating table in Sec. 7.7. (3) The current design can detect either texture or material with different labeling, but does not support simultaneous detection of both. Fig. 20d in Sec. 7.7 evaluates a scenario containing eight surfaces - spanning *four materials* and *four surfaces of different roughness levels within one same wooden material* – treated as eight distinct labels. This result demonstrates the potential of extending mmTexora to simultaneously classify two properties: material type and surface roughness.

Future Work and Applications: In this paper, we deployed mmTexora using commercial mmWave front-ends and a robotic arm for evaluation, demonstrating its potential for integration into more robust and adaptive robotic perception systems. Our observations from Fig. 6a (Radar itself is under vibrations from the near-field phase responses) and Sec. 7.7 show the potential of mmTexora to operate under robot dynamics in realistic confessions, where the radar experienced vibrations as the arm moved across the operating platform. In addition, our approach could enable mobile robots to manipulate conventionally hard-to-grasp objects for practical real-world applications. We believe several limitations of mmTexora that can be improved in future work: (1) supporting simultaneous material and roughness sensing; (2) supporting liquid and elastic objects; (3) supporting operation under extreme clutter, such as outdoors and scenes with dynamics and crowds.

9 Conclusion

This paper presents mmTexora, a 77-GHz mmWave radar-based system for the detection of fine-grained surface materials and roughness. We show the feasibility of leveraging micro-motion induced by ambient vibrating sources, like ventilation and machinery, to extract the embedded surface scattering information related to the surface’s properties in a static setup. Experimental results demonstrate the capability of mmTexora to achieve tens of micrometer granularity in surface sensing. Through extensive experiments on a large set of different surface materials, we demonstrate submillimeter-level roughness detection in a fixed location.

Acknowledgments

We acknowledge support from the NSF (2106921, 2433903, 1942902, 2111751), ONR (N000142412410, N000142412062), DARPA,

AFRETEC, CMU Block Center, CMU MFI, CMU Mill 19, and CyLab-Enterprise. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the above.

References

- [1] 2024. AWR1843BOOST. "https://www.ti.com/tool/AWR1843BOOST".
- [2] 2024. DCA1000EVM. "https://www.ti.com/tool/DCA1000EVM".
- [3] 2024. Precision Devices, Inc. Surface roughness terminology and parameters. https://www.predev.com/pdffiles/surface_roughness_terminology_and_parameters.pdf
- [4] 2025. Motoman GP4 High-Speed Industrial Robots. "https://www.motoman.com/en-us/products/robots/industrial/assembly-handling/gp-series/gp4".
- [5] 2025. OPTICAL PROFILOMETER VR 5000. "https://www.keyence.com/landing/measure-sys/lp_wide-area-measurement_01084784.jsp?aw=google-kaenVR323207ee-br&gad_source=1&gad_campaignid=653401887&gbraid=0AAAAADJ0X7iFCeEq6xJydGHpGRzZ72CN&gclid=Cj0KCQjwjo7DBhCrARIsACWauSkYmWTe0N6cnihrnAS5uzVM6QIOdqr61-D-7Suo965SVVpBjT_bb0aAi4GELw_wcB".
- [6] Waleed H Abdulla, David Chow, and Gary Sin. 2003. Cross-words reference template for DTW-based speech recognition systems. In *TENCON 2003. Conference on convergent technologies for Asia-Pacific region*, Vol. 4. IEEE, 1576–1579.
- [7] Stefano Aldini, Avinash K Singh, Daniel Leong, Yu-Kai Wang, Marc G Carmichael, Dikai Liu, and Chin-Teng Lin. 2022. Detection and estimation of cognitive conflict during physical human–robot collaboration. *IEEE Transactions on Cognitive and Developmental Systems* 15, 2 (2022), 959–968.
- [8] Hélio Azevedo, Roseli AF Romero, and José Pedro Ribeiro Belo. 2017. Reducing the gap between cognitive and robotic systems. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 1049–1054.
- [9] Daniele Barchiesi, Dimitrios Giannoulis, Dan Stowell, and Mark D. Plumbley. 2015. Acoustic Scene Classification: Classifying environments from the sounds they produce. *IEEE Signal Processing Magazine* 32, 3 (2015), 16–34. doi:10.1109/MSP.2014.2326181
- [10] Petr Beckmann and Andre Spizzichino. 1987 - 1963. *The scattering of electromagnetic waves from rough surfaces*. Artech House, Norwood, MA.
- [11] Roger Bostelman and James Albus. 2008. Sensor experiments to facilitate robot use in assistive environments. In *Proceedings of the 1st international conference on Pervasive Technologies Related to Assistive Environments*. 1–9.
- [12] Xingyu Cai, Tingyang Xu, Jinfeng Yi, Junzhou Huang, and Sanguthevar Rajasekaran. 2019. *DTWNet: a dynamic time warping network*. Curran Associates Inc., Red Hook, NY, USA.
- [13] Roberto Calandra, Serena Ivaldi, Marc Peter Deisenroth, and Peters. 2015. Learning torque control in presence of contacts using tactile sensing from robot skin. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. 690–695. doi:10.1109/HUMANOIDS.2015.7363429
- [14] Johann Fauzi. 2022. Time series classification: A review of algorithms and implementations. *Machine Learning (Emerging Trends and Applications)* (2022).
- [15] A.K. Fung and K.S. Chen. 2004. An update on the IEM surface backscattering model. *IEEE Geoscience and Remote Sensing Letters* 1, 2 (2004), 75–77. doi:10.1109/LGRS.2004.826564
- [16] Adrian K. Fung. 1994. *Microwave scattering and emission models and their applications*. Artech House.
- [17] Kirk W Gossage, Tomasz S Tkaczyk, Jeffrey J Rodriguez, and Jennifer K Barton. 2003. Texture analysis of optical coherence tomography images: feasibility for tissue classification. *Journal of biomedical optics* 8, 3 (2003), 570–575.
- [18] Jun Han, Albert Jin Chung, Manal Kumar Sinha, Madhumitha Harishankar, Shijia Pan, Hae Young Noh, Pei Zhang, and Patrick Tague. 2018. Do You Feel What I Hear? Enabling Autonomous IoT Device Pairing Using Different Sensor Types. In *2018 IEEE Symposium on Security and Privacy (SP)*. 836–852. doi:10.1109/SP.2018.00041
- [19] Ali Ismail-Fawaz, Maxime Devanne, Jonathan Weber, and Germain Forestier. 2022. Deep learning for time series classification using new hand-crafted convolution filters. In *2022 IEEE International Conference on Big Data (Big Data)*. 972–981.
- [20] Hassan Ismail Fawaz, Benjamin Lucas, Germain Forestier, Charlotte Pelletier, Daniel F Schmidt, Jonathan Weber, Geoffrey I Webb, Lhassane Idoumghar, Pierre-Alain Muller, and François Petitjean. 2020. Inceptiontime: Finding alexnet for time series classification. *Data Mining and Knowledge Discovery* 34, 6 (2020), 1936–1962.
- [21] Zhaohui Li, Wei Luo, Yongmin Zhang, Jianxi Chen, Yuanhao Shu, and Yaoxue Zhang. 2024. ASLiquid: Non-Intrusive Liquid Counterfeit Identification with Your Earphones. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems (Hangzhou, China) (SenSys '24)*. Association for Computing Machinery, New York, NY, USA, 41–53. doi:10.1145/3666025.3699321
- [22] Tangyou Liu, Tinghua Zhang, Jay Katupitiya, Jiaole Wang, and Liao Wu. 2024. Haptics-Enabled Forceps With Multimodal Force Sensing: Toward Task-Autonomous Surgery. *IEEE/ASME Transactions on Mechatronics* 29, 3 (2024), 2208–2219. doi:10.1109/TMECH.2023.3329564
- [23] Yawen Liu, Shihan Lu, and Heather Culbertson. 2022. Texture Classification by Audio-Tactile Crossmodal Congruence. In *2022 IEEE Haptics Symposium (HAPTICS)*. 1–7. doi:10.1109/HAPTICS52432.2022.9765614
- [24] Shan Luo, Wenzhen Yuan, Edward Adelson, Anthony G. Cohn, and Raul Fuent. 2018. ViTac: Feature Sharing Between Vision and Tactile Sensing for Cloth Texture Recognition. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. 2722–2727. doi:10.1109/ICRA.2018.8460494
- [25] Rajni V Patel, S Farokh Atashzar, and Mahdi Tavakoli. 2022. Haptic feedback and force-based teleoperation in surgical robotics. *Proc. IEEE* 110, 7 (2022), 1012–1027.
- [26] Akarsh Prabhakara, Vaibhav Singh, Swaran Kumar, and Anthony Rowe. 2020. Osprey: a mmWave approach to tire wear sensing. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services (Toronto, Ontario, Canada) (MobiSys '20)*. Association for Computing Machinery, New York, NY, USA, 28–41. doi:10.1145/3386901.3389031
- [27] Christopher Reardon, Jason M Gregory, Kerstin S Haring, Benjamin Dossett, Ori Miller, and Aniekang Inyang. 2024. Augmented Reality Visualization of Autonomous Mobile Robot Change Detection in Uninstrumented Environments. *ACM Transactions on Human-Robot Interaction* 13, 3 (2024), 1–30.
- [28] Joseph M. Romano and Katherine J. Kuchenbecker. 2014. Methods for robotic tool-mediated haptic surface recognition. In *2014 IEEE Haptics Symposium (HAPTICS)*. 49–56. doi:10.1109/HAPTICS.2014.6775432
- [29] Hailan Shanbhag, Sohrab Madani, Akhil Isanaka, Deepak Nair, Saurabh Gupta, and Haitham Hassanieh. 2023. Contactless Material Identification with Millimeter Wave Vibrometry. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services (Helsinki, Finland) (MobiSys '23)*. Association for Computing Machinery, New York, NY, USA, 475–488. doi:10.1145/3581791.3596850
- [30] Ruiyi Shen and Yasaman Ghasempour. 2023. Scattering from Rough Surfaces in 100+ GHz Wireless Mobile Networks: From Theory to Experiments. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking (Madrid, Spain) (ACM MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 92, 15 pages. doi:10.1145/3570361.3613306
- [31] Matti Strese, Lara Brudermueller, Jonas Kirsch, and Eckehard Steinbach. 2020. Haptic Material Analysis and Classification Inspired by Human Exploratory Procedures. *IEEE Transactions on Haptics* 13, 2 (2020), 404–424. doi:10.1109/TOH.2019.2952118
- [32] Matti Strese, Clemens Schuwerk, Albert Iepure, and Eckehard Steinbach. 2017. Multimodal Feature-Based Surface Material Classification. *EEE Trans. Haptics* 10, 2 (April 2017), 226–239. doi:10.1109/TOH.2016.2625787
- [33] Matti Strese, Clemens Schuwerk, Albert Iepure, and Eckehard Steinbach. 2017. Multimodal Feature-Based Surface Material Classification. *IEEE Transactions on Haptics* 10, 2 (2017), 226–239. doi:10.1109/TOH.2016.2625787
- [34] Yizhou Tan, Haojun Ai, Shengchen Li, and Mark D. Plumbley. 2024. Acoustic Scene Classification Across Cities and Devices via Feature Disentanglement. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.* 32 (Jan. 2024), 1286–1297. doi:10.1109/TASLP.2024.3353578
- [35] Tasbolat Taunyazov, Luar Shui Song, Eugene Lim, Hian Hian See, David Lee, Benjamin C.K. Tee, and Harold Soh. 2021. Extended Tactile Perception: Vibration Sensing through Tools and Grasped Objects. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Prague, Czech Republic). IEEE Press, 1755–1762. doi:10.1109/IROS51168.2021.9636677
- [36] Stephen Tian, Frederik Ebert, Dinesh Jayaraman, Mayur Mudigonda, Chelsea Finn, Roberto Calandra, and Sergey Levine. 2019. Manipulation by feel: Touch-based control with deep predictive models. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 818–824.
- [37] Thakshanth Uthayakumar, Mohamed Rouzin Azar Mohamed Rifa, and Jayathu G Samarawickrama. 2020. Teleoperation of a Robotic Arm through Tactile Sensing, Visual and Haptic Feedback. In *2020 5th International Conference on Control and Robotics Engineering (ICCRE)*. IEEE, 63–67.
- [38] Mary Catherine VG, Binu Paul, et al. 2022. Wearable Fabric Tactile Sensors for Robotic Elderly Assistance. In *2022 Smart Technologies, Communication and Robotics (STCR)*. IEEE, 1–5.
- [39] M. Visscher and K.G. Struik. 1994. Optical profilometry and its application to mechanically inaccessible surfaces Part I: Principles of focus error detection. *Precision Engineering* 16, 3 (1994), 192–198. doi:10.1016/0141-6359(94)90124-4
- [40] Ju Wang, Jie Xiong, Xiaojiang Chen, Hongbo Jiang, Rajesh Krishna Balan, and Dingyi Fang. 2017. TagScan: Simultaneous Target Imaging and Material Identification with Commodity RFID Devices. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking (Snowbird, Utah, USA) (MobiCom '17)*. Association for Computing Machinery, New York, NY, USA, 288–300. doi:10.1145/3117811.3117830
- [41] Jonas Weiß and Avik Santra. 2018. One-Shot Learning for Robust Material Classification Using Millimeter-Wave Radar System. *IEEE Sensors Letters* 2, 4 (2018), 1–4. doi:10.1109/LESEN.2018.2878041

- [42] Chenshu Wu, Feng Zhang, Beibei Wang, and K. J. Ray Liu. 2020. mSense: Towards Mobile Material Sensing with a Single Millimeter-Wave Radio. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 4, 3, Article 106 (Sept. 2020), 20 pages. doi:10.1145/3411822
- [43] Binbin Xie, Jie Xiong, Xiaojiang Chen, Eugene Chai, Liyao Li, Zhanyong Tang, and Dingyi Fang. 2019. Tagtag: material sensing with commodity RFID. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems* (New York, New York) (*SenSys '19*). Association for Computing Machinery, New York, NY, USA, 338–350. doi:10.1145/3356250.3360027
- [44] Han Xu, Mingqi Chen, Gaofeng Li, Lei Wei, Shichi Peng, Haoliang Xu, and Qiang Li. 2025. An Immersive Virtual Reality Bimanual Telerobotic System With Haptic Feedback. *arXiv preprint arXiv:2501.00822* (2025).
- [45] Dawei Yan, Panlong Yang, Fei Shang, Weiwei Jiang, and Xiang-Yang Li. 2024. Wi-Painter: Fine-grained Material Identification and Image Delineation Using COTS WiFi Devices. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 4, Article 203 (Jan. 2024), 25 pages. doi:10.1145/3633809
- [46] Vahid Yazdian, Ruiyi Shen, and Yasaman Ghasempour. 2025. *RoboTera: Non-Contact Friction Sensing for Robotic Grasping via Wireless Sub-Terahertz Perception*. Association for Computing Machinery, New York, NY, USA, 172–185. <https://doi.org/10.1145/3715014.3722072>
- [47] 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* (Tokyo, Japan) (*UIST '16*). Association for Computing Machinery, New York, NY, USA, 833–841. doi:10.1145/2984511.2984515
- [48] Diana Zhang, Jingxian Wang, Junsu Jang, Junbo Zhang, and Swarun Kumar. 2019. On the Feasibility of Wi-Fi Based Material Sensing. In *The 25th Annual International Conference on Mobile Computing and Networking* (Los Cabos, Mexico) (*MobiCom '19*). Association for Computing Machinery, New York, NY, USA, Article 41, 16 pages. doi:10.1145/3300061.3345442
- [49] Zhan Zhang, Denghui Song, Anfu Zhou, and Huadong Ma. 2024. airTac: A Contactless Digital Tactile Receptor for Detecting Material and Roughness via Terahertz Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8, 3, Article 141 (Sept. 2024), 37 pages. doi:10.1145/3678586
- [50] Tianyue Zheng, Zhe Chen, Jun Luo, Lin Ke, Chaoyang Zhao, and Yaowen Yang. 2021. SiWa: see into walls via deep UWB radar. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking* (New Orleans, Louisiana) (*MobiCom '21*). Association for Computing Machinery, New York, NY, USA, 323–336. doi:10.1145/3447993.3483258