

On the Feasibility of Wi-Fi Based Material Sensing

Diana Zhang
Carnegie Mellon University
dianaz1@andrew.cmu.edu

Jingxian Wang
Carnegie Mellon University
jingxian@cmu.edu

Junsu Jang
Massachusetts Institute of
Technology
junsuj@mit.edu

Junbo Zhang
Carnegie Mellon University
junboz2@andrew.cmu.edu

Swarun Kumar
Carnegie Mellon University
swarun@cmu.edu

ABSTRACT

Wireless sensing has demonstrated the potential of using Wi-Fi signals to track people and objects, even behind walls. Yet, prior work in this space aims to merely detect the presence of objects around corners, rather than their type. In this paper, we explore the feasibility of the following research question: “Can commodity Wi-Fi radios detect both the location and *type* of moving objects around them?”. We present IntuWition, a complementary sensing system that can sense the location and type of material of objects in the environment, including those out of line-of-sight. It achieves this by sensing wireless signals reflected off surrounding objects using commodity Wi-Fi radios, whose signals penetrate walls and occlusions. At the core of IntuWition is the idea that different materials reflect and scatter polarized waves in different ways. We build upon ideas from RADAR Polarimetry to detect the material of objects across spatial locations, despite mobility of the sensing device and the hardware non-idealities of commodity Wi-Fi radios. A detailed feasibility study reveals an average accuracy of 95% in line-of-sight and 92% in non-line-of-sight in classifying five types of materials: copper, aluminum, plywood, birch, and human. Finally, we present a proof-of-concept application of our system on an autonomous UAV that uses its onboard Wi-Fi radios to sense whether an occlusion is a person versus another UAV.

ACM Reference Format:

Diana Zhang, Jingxian Wang, Junsu Jang, Junbo Zhang, and Swarun Kumar. 2019. On the Feasibility of Wi-Fi Based Material Sensing. In

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom '19, October 21–25, 2019, Los Cabos, Mexico

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6169-9/19/10...\$15.00

<https://doi.org/10.1145/3300061.3345442>

The 25th Annual International Conference on Mobile Computing and Networking (MobiCom '19), October 21–25, 2019, Los Cabos, Mexico. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3300061.3345442>

1 INTRODUCTION

Recent years have seen an explosion in wireless sensing and tracking research, from people to emotions to objects. However, most existing work in the wireless sensing space do not sense the precise material composition of the objects they track. Yet, in many occasions wireless radios may need to be aware of the type of objects in the environment, beyond their location. For instance, consider autonomous UAVs that may choose to use their onboard Wi-Fi radios to detect if an object around the corner is a person or another drone. This could allow it to respond differently in either case – for example by giving the human a larger berth versus moving up vertically to avoid the drone (see Fig. 1). Beyond UAVs, object type sensing could transform ordinary Wi-Fi devices into sensors that identify objects around the corner, with applications in search-and-rescue, smart occupancy sensing, vehicular safety and beyond.

Unfortunately, state-of-the-art sensing solutions fall short of localizing and discerning occluded obstacles. For instance, cameras can both localize and identify objects, but operate solely in line-of-sight. This could be addressed with high-resolution aircraft RADARs; however, these bulky devices are not portable and are highly expensive [4]. Further, newer radars designed for flight are often at high frequencies for higher resolution, but this comes at the cost of reduced ability to propagate through and around materials and 6x the power consumption over a typical Wi-Fi router [3]. Wi-Fi based sensing systems [5, 34] can penetrate walls and obstacles unlike light to track hidden objects, but have not been comprehensively explored for material sensing. In this paper, we ask: “Can commodity Wi-Fi radios sense the location and *type* of moving objects?”

We present IntuWition, a system that explores the feasibility of identifying both the material and location of surrounding objects using commodity Wi-Fi. We envision IntuWition

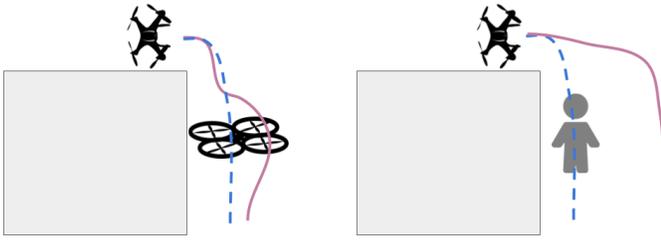


Figure 1: IntuWition helps detect the type of material of an object hidden from view. This could assist, say a drone’s path planning algorithm, to change its planned path (blue) to a new path (purple). For example, it could swerve upwards to avoid another drone, but would give a human a wide berth.

as a complementary sensing modality to add non-line-of-sight and materials-sensing capabilities to existing sensing systems. We deploy Wi-Fi radios indoors and outdoors on a college campus to explore the materials sensing capabilities of WiFi. IntuWition processes wireless channels between pairs of commodity Wi-Fi radios to accurately estimate and track the location and material of different objects in its surroundings, including wooden, metal, and human obstacles. We further present a proof-of-concept application demonstrating IntuWition’s performance on autonomous UAVs in distinguishing between occluded humans vs. other drones both indoors and outdoors.

At the heart of IntuWition is an approach that infers material properties of objects in the environment by measuring the wireless signals that reflect off them. In particular, it captures a specific property: the *polarization* of the reflected waves from an object. When a polarized incident wireless signal reflects off the surface of different objects, it behaves differently based on the texture and material of the objects. Specifically, metal absorbs and re-emits the wave, resulting in a change of polarization pattern; smooth surfaces reflect the waves with polarization intact; and rough surfaces scatter the waves and diffuse polarization. In effect, this causes a change in the observed polarization of the reflected wave. We demonstrate that by measuring the polarization of reflected waves from surrounding objects, one can infer the material they are composed of. We draw inspiration from the remote sensing community, where aircrafts image the topography using RADAR polarimetry [33, 42] to distinguish between trees, open fields and even infer soil moisture using reflected signals.

To bring RADAR polarimetry to commodity Wi-Fi radios, IntuWition measures the Wi-Fi signal from a vertically-polarized transmitter antenna to three mutually-perpendicular polarized receiving antennas, all of which are atop a Wi-Fi enabled device. We then compare the measured power of the

signal across these antennas to infer its material composition. IntuWition builds a multi-layer perceptron model to achieve this independent of other parameters that influence received signals such as the size, thickness, texture, distance, and angle of an object. A key challenge in IntuWition’s design is the ability to detect and discard signals that are reflected and scattered by multiple objects sequentially and therefore experience a hybrid change in polarization that corresponds to none of their underlying materials. IntuWition achieves this by reverse-engineering and recognizing the unique ways in which a Wi-Fi signal’s polarization compounds when it bounces off many objects (see Sec. 5).

A second challenge IntuWition addresses is to separate the signals that reflect exclusively off each individual object in the environment in 3-D space, prior to analyzing its material composition. IntuWition does so by processing these signals to measure the time-of-flight they experience as they reflect off each obstacle over brief slices of time. Additionally, mobility of the sensing autonomous systems can allow for views the same object from multiple perspectives. These samples further allow us to triangulate its 3-D physical position and to filter out the object’s reflected signal for material classification. A key challenge in doing so pertains to the hardware non-idealities of commodity wireless radios that induce random errors in the measured time-of-flight owing to arbitrary timing offsets between the Wi-Fi transmitter and receiver, both of which are two distinct radio chips with different clocks. IntuWition addresses this challenge by simultaneously sending signals from the transmitter along two pathways: a wired pathway and a wireless pathway, only the latter of which is impacted by the environment. We then estimate timing offsets from the wired pathway to correct for its effect on the wireless pathway. The rest of this paper describes how IntuWition addresses this and other hardware non-idealities such as carrier frequency offsets, sampling offsets, and phase shifts between the transmitting and receiving RF chains. Further, we discuss the limitations of IntuWition in Sec. 7: It often misses small, fast-moving or well-shielded objects, and cannot tell apart materials with similar polarization characteristics (e.g. two different humans).

We implement IntuWition on Intel Galileo boards equipped with commodity Intel 5300 Wi-Fi cards [16]. We perform a detailed feasibility study in large indoor and outdoor spaces in a university campus and distinguish between common material types: copper, aluminum, humans, plywood and birch, in a variety of line-of-sight, non-line-of-sight, stationary, and mobile settings. We further mount our Wi-Fi platform on a UAV (DJI Matrice 100) to classify between humans and drones for indoor and outdoor testing. Our results reveal the following:

- Our system classifies between 5 types of material (copper, aluminum, plywood, birch, and human) of a variety of sizes and orientations at an accuracy of 95% in line-of-sight and 92% in non-line-of-sight settings.
- Our system can classify objects that are 0.42 m apart in line-of-sight and 0.55 m apart in non-line-of-sight.
- Our system classifies between humans and drones with accuracy averaging 89% at UAV speeds of up to 2 m/s in dynamic indoor and outdoor settings.

Contributions: This paper presents IntuWition, to our knowledge, the first system that explores the use of commodity Wi-Fi radios in discerning objects in the environment using polarization. We demonstrate how IntuWition detects the material and location of occluded obstacles, opening up applications for autonomous UAVs and beyond. IntuWition is fully implemented and evaluated indoors and outdoors.

2 RELATED WORK

Related work falls under three broad categories:

Wireless Sensing: Recent years have seen extensive work in sensing the environment using wireless communication radios. Past work has leveraged variations in Wi-Fi signals for activity recognition [34, 47], occupancy sensing [6, 11, 12], imaging [20, 37] and location tracking [5, 22, 25, 49, 50]. Indeed, a range of device-free applications have been proposed by fine-grained motion tracking of users, sensitive enough for even breath/heart rate tracking [52] and keystroke sensing [7]. Perhaps most closely related to our work are recent material detection systems [13, 46, 48] that can detect material by sensing how signals are attenuated by objects placed between the transmitter and receiver. Much of this work assumes multiple access points or bulky infrastructure in the environment and are therefore ill-suited for mobile platforms such as UAVs.

This paper instead seeks to discover the location and material of multiple potentially occluded objects in the environment using moving Wi-Fi radios on a mobile platform without supporting infrastructure. Key to our approach is the use of polarization of reflected signals to detect the material of objects using commodity Wi-Fi radios.

RADAR & Polarimetry: RADAR systems use synthetic aperture radar polarimetry for topography imaging [9, 33, 38, 39]. Aircraft and satellites can even use polarization of received RADAR signals to measure soil moisture of farmland [42]. Recent years has seen the persistence of polarimetry largely unchanged as a sensing modality but used with new materials [29] and for new applications [54]. In contrast to this rich literature, our approach seeks to bring radar polarimetry to light-weight commodity Wi-Fi cards. For this, we overcome many challenges such as the near-far effect

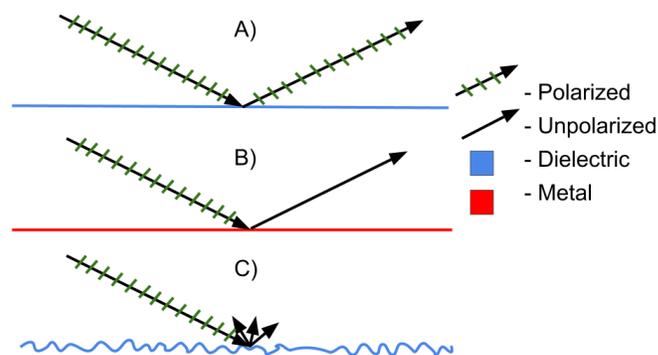


Figure 2: Scattering off (A) a smooth surface, (B) a smooth metallic surface, and (C) a rough surface produces differently polarized signals.

and imperfections of such radios. We further distinguish ourselves from existing polarimetric systems in that our proximity is much closer than an airplane or a satellite, where polarimetry systems are commonly deployed.

Project Soli, Google’s RADAR-based system for object material recognition [53] assumes extremely close proximity between the object and the RADAR system. Synthetic Aperture Radar has also been considered for device tracking of Wi-Fi [26], cellular [27] and RFID [41, 45] tags. IntuWition additionally uses polarization for device-free material sensing of surrounding objects.

Other Sensing Systems for UAVs and Beyond: Cameras [31], LIDARs [36] and depth imaging [19] are often used to detect occluded objects, for instance by autonomous vehicles and UAVs. While these sense obstacles in direct field of view, they fail to detect objects that are occluded (e.g. hidden pedestrian or drones around the corner) – a challenge IntuWition targets. IntuWition aims to do so without supporting infrastructure in contrast to [14, 28, 35]. We further seek to sense the environment, not characteristics of a UAV as in [32]. Finally, we distinguish ourselves from [23], likely the most related paper to our work, in that we seek to sense the environment via reflections rather than penetrations.

Proximity RADARs have been used to detect collisions in automobiles [15] and UAVs [1], although they do not sense materials. Recent work has explored mm-wave radars to sense the immediate environment of UAVs and autonomous cars [21, 56]. In contrast, we leverage commodity Wi-Fi for material sensing for two reasons: (1) We can detect occluded objects, unlike mm-wave signals that are blocked by walls; (2) Wi-Fi is inexpensive, light-weight and already present in many mobile platforms, such as phones, tablets and drones.

3 POLARIMETRY FOR SENSING

This section provides a brief primer on how capturing the polarization of signals that scatter off an object provides information about the material it is made of. Fig. 2 illustrates three extreme instances of scattering responses for a perfectly polarized incident wave: (A) The incident wave reflects off a perfectly smooth non-metallic surface, maintaining linear polarization. (B) The incident wave reflects off a perfectly smooth metallic surface, eliminating linear polarization [24]. (C) The incident wave scatters off a perfectly rough surface, which may or may induce polarization based on the material. Hence, by measuring the polarization and power of the received waves, we can distinguish objects based on texture and material (metallic or otherwise).

To better understand this phenomenon, we consider a 2-D case for simplicity, a surface defined by parallel and perpendicular vectors p and s , respectively. As described in [17], an incident s -polarized signal can be denoted as $E_s(t) = Ae^{j\omega t} \hat{e}^s$. Upon reflection off a surface, amplitude, phase, and orientation can all be affected, resulting in a reflected signal $E_r(t) = Ae^{j\omega t} (r_{ss} \hat{e}^s + r_{sp} \hat{e}^p)$, where r_{ss} and r_{sp} are complex numbers describing how the s -oriented incident wave has changed in phase and power along each of the s and p components. Note that depending on the two phase value shifts, the locus of our signal might change from linear polarization to circular or elliptical – if they change by the same amount, or by a factor of π , linearly polarization is preserved.

Because of perpendicular linearly-polarized receiving antennas, IntuWition measures \hat{e}^s and \hat{e}^p separately, thus being able to measure material-specific parameters r_{ss} and r_{sp} . Having antennas that allow us to receive different polarizations of data opens the door to polarimetry. Furthermore, receiving these signals separately allows us to take the ratio of these signals to remove power and phase shifts due to propagation, which affects signals similarly regardless of their orientation.

We note that in practice this property is influenced by the location, geometry, orientation and texture of reflecting surfaces in the environment besides material composition. The following sections present models to process wireless channels to decouple the reflector’s material composition from all the remaining quantities.

4 INTUWITON IN A NUTSHELL

IntuWition’s objective is to allow a Wi-Fi enabled mobile devices to obtain both the 3-D location and the material composition of surrounding objects, including those occluded from view. It aims to do this using existing commodity Wi-Fi radios on the drone, without requiring supporting wireless infrastructure in the environment. In other words, an

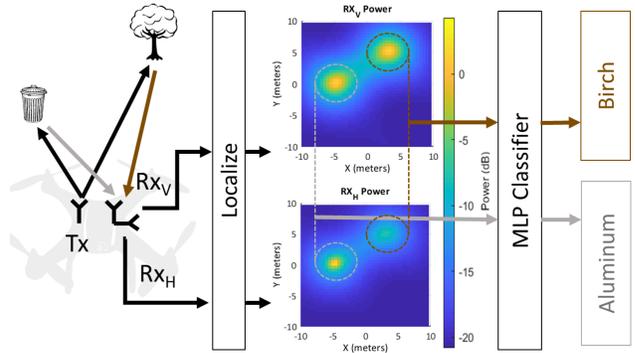


Figure 3: IntuWition Workflow: An example in 2-D (for simplicity) shows how IntuWition isolates locations of two objects and compares powers across horizontally and vertically polarized antennas to infer material type.

IntuWition device must both transmit and receive Wi-Fi signals and analyze them to report the location and material of surrounding reflectors. Fig. 3 presents an illustrative workflow of IntuWition.

4.1 Learning Material Composition

IntuWition determines the material composition of any object-of-interest in the environment by studying the polarization of the signal reflected from it. Specifically, IntuWition first uses the known time-of-flight of the reflected signal from an object to develop a filter that outputs the wireless channel component that corresponds solely to the reflected path. It then processes this signal component on a horizontal and vertically polarized receiving antenna and measures their ratio to determine the extent to which the reflection impacted the polarization of the transmitted signal. IntuWition measures this ratio across a range of frequency bands, developing a vector that captures the impact of the material on signal polarization. We then build a model based on multi-layer perceptrons to classify the material and type of the object-of-interest. We also compare our model with five other typical machine learning models.

Extracting Features: Among the key challenges in material classification is identifying the best features that are unique to materials. IntuWition relies on the change in polarization of the received signal from the transmitted signal across different frequencies. Specifically, we filter different *subsets* of the wireless channels, each spanning a different range of frequencies to obtain the wireless channel component reflected from the object-of-interest. We then compare these channels at the horizontal and vertical antennas to study their polarization across subsets of frequency bands. Finally, we feed the ratio between channels at the horizontal and vertical antennas into our model. IntuWition leverages

a Multi-Layer Perceptron model to process this feature vector and infer the object material and type, while remaining robust to device movement and signal multipath. Sec. 5.1 details IntuWition’s learning models.

Sensing Mobile Objects over Time: Perhaps the most challenging objects to track are objects that are themselves moving with time. For instance, consider two objects whose trajectories intersect. While IntuWition can identify an approximate relative location of each moving object along their respective trajectories, it may erroneously swap the trajectories of the two when they intersect. Sensing materials of objects provides a natural solution to this problem. Specifically, IntuWition can use the polarization of reflected signals from each object to identify it and accurately track its position over time without ambiguity. Sec. 5.2 describes how we exploit the synergy between localization and material sensing to improve each other’s performance.

4.2 Locating Surrounding Objects

To accurately locate surrounding objects, IntuWition analyzes the wireless channel-state-information available at a compact commodity 3-antenna Wi-Fi receiver from a co-located Wi-Fi transmitter on the drones. These channels are a combination of signals propagating along different paths as they reflect off various surrounding objects, as well as the direct path between the transmitter and receiver. IntuWition processes these wireless channels across Wi-Fi frequency bands and measures the time-of-flight experienced by the signals along each path. When multiplied by the speed of light, this provides the distance traversed by the signal from the transmitter to the reflector and then back to the receiver. By computing these distances from multiple perspectives as the drone flies in the 3-D space, IntuWition triangulates the 3-D position of the object-of-interest.

While there has been much work on localizing surrounding objects using RADAR [15, 44] and recent work on wireless material sensing applications [13, 46, 48], bringing it to inexpensive and commodity Wi-Fi radios on mobile systems without external supporting infrastructure brings forth several new challenges that IntuWition must address:

Disentangling Signal Paths: First, IntuWition must separate signal paths from various objects surrounding the mobile device. IntuWition achieves this by actively exploiting the mobility of the device itself (e.g. a moving UAV or tablet). Sec. 6.1 shows how by processing the received wireless channels across spatial locations and frequency of operation of Wi-Fi radios, one can simultaneously obtain time-of-flight and angle-of-arrival of surrounding objects.

Locating Occluded Objects: Second, IntuWition must develop algorithms that can analyze weak reflections from

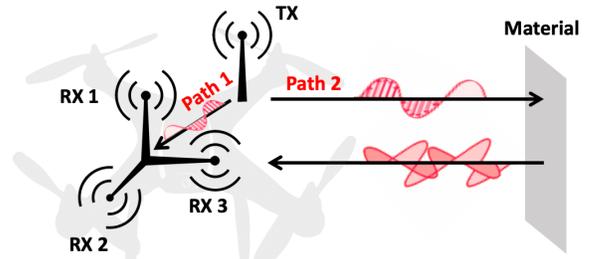


Figure 4: IntuWition transmits signals from a vertically polarized antenna to 3 mutually perpendicular receive antennas and processes polarization changes due to reflections from materials.

distant objects and those that are occluded by other objects (e.g. walls, trees, etc.). IntuWition achieves this by sensing changes in the wireless channels from weak reflections to detect the presence of the moving objects that produce them. A key challenge in implementing this is to filter out the movement of the drone itself, as it also changes wireless channels, over and above the movement of the object-of-interest. IntuWition therefore develops a novel background subtraction algorithm that eliminates the effect of the drone’s movement in the wireless channels to detect the movement of surrounding objects. Sec. 6.2 describes our approach in greater detail.

5 MATERIAL SENSING

This section describes how IntuWition senses the material properties of a surrounding object. We assume the received signal has been pre-processed to localize and isolate the signal component from a specific object of interest (detailed in Sec. 6). We describe how we bring RADAR polarimetry to the Wi-Fi context, while mitigating the limitations of commodity Wi-Fi radios.

Hardware Setup: IntuWition’s setup consists of four Wi-Fi radios – a transmitting device broadcasts packets from a vertically polarized antenna to three receiving light-weight Wi-Fi chips, each with a mutually perpendicularly polarized receiver antenna on port one. Note that inexpensive omnidirectional whip antennas can be used as vertically polarized antennas because of their natural propagation pattern. We connect a wire to the transmitter and, via a splitter and attenuated cable, connect this to antenna port 2 on each receiving chip. This provides resilience to phase errors from hardware offsets (details in Sec. 6.1).

Fig. 4 shows an illustration of how the polarization pattern of signal is changed after reflections from a material. If no reflector is present in the environment, the three mutually perpendicular receiver antennas RX1, RX2, RX3 will only receive the single signal polarization pattern of the direct path (Path1). In the presence of a reflector, the receive antennas

also obtain signals along *Path2* whose pattern of polarization depends on the material of the reflector.

IntuWition uses three receiving antennas instead of two as is used in the traditional RADAR polarimetry because our transmitter and receiver are at about the same elevation level of the obstacle we are sensing. As a result, modeling the sensed object as a 2-D reflecting plate (as is the case from a satellite or a high plane) is no longer valid. Even when movement is along four controllable degrees of freedom (say on a quadcopter drone), three antennas are sufficient to collect information about scattering from various surrounding objects (including the floor), despite the movement of the platform itself.

Approach: IntuWition employs Multi Layer Perceptrons to classify different materials for a given object. Our system takes as input the polarization of the reflected signal from a given object across frequencies. It then extracts features via a multi-layer perceptron that closely correlate with material composition of the object. A key challenge we face in designing our material classification algorithm is to decouple various other factors that influence the wireless channels, such as object size, texture and orientation. The rest of this section details our solutions to deal with these challenges as well as our choice of features and classification algorithm.

5.1 Building the Machine Learning Model

Prior to extracting features, IntuWition uses its object localization algorithm (details in Sec. 6) to filter the wireless channel from an object of interest. In particular, it measures this channel at all receive antennas across OFDM subcarriers. Further, IntuWition repeats the localization algorithm over many different subsets of Wi-Fi operating frequencies to obtain the wireless channel component reflected off the object. In doing so, IntuWition retrieves a vector of wireless channels per-antenna from the object versus frequency – both across different subsets of frequency bands and subcarriers.

Next, IntuWition computes the ratios of the channel received between pairs of antennas (Fig. 4). By doing so, we eliminate the effect of distance of the object on signal power and focus instead on polarization, which is highly correlated with the material composition of the object. The resulting three channel ratios sampled across frequencies form a vector of over a thousand elements. Feeding all of these as features into our machine learning model would be counter-productive, given that it requires a complex higher-dimensional classification model that is highly susceptible to over-fitting.

However, we find that channel ratios sampled across all frequencies are not all equally important as features, owing to differences in the quality of channels (e.g. channel fading and interference on some frequencies) or the number of measurements available. Hence, IntuWition selects

a subset of features from the wireless channel ratios across frequencies. We use a greedy algorithm to rank features according to number of successful measurements across the whole channel and select the top 200 ones. Then we remove the samples which have missing values among those 200 features. In doing so, we can reduce the number of features to avoid over-fitting.

While a range of machine learning classification models are available, IntuWition employs Multi-Layer Perceptrons (MLPs) to train its classifier based on a theoretical analysis and an empirical comparison. MLPs have been widely observed to generalize well compared to hand-crafted features in recent deep learning literature [43], since they automatically learn features from input data. We construct two versions of MLPs which take raw feature values as input and outputs either the material class or an object type. The latter pertains to our proof-of-concept demonstration of our system on a UAV, which classifies between two classes of objects – drones and humans who may have multiple (yet different) classes of materials on their surface.

Besides MLP, we considered alternative machine learning models. Sec. 9 presents an empirical comparison of various machine learning models: MLP, RBF-SVM, k -NN, PCA, and Naïve Bayes, and shows that MLPs provide the maximum classification accuracy for us.

Training and Testing: We train the model by feeding the top 200 features to Multi-Layer Perceptrons. These features were selected as the most impactful out of all the power ratios between pairs of antennas captured at frequencies across large bandwidth. To avoid the overfitting of the training model, we collect the training data in different radio environments (indoor and outdoor), across multiple days. We collect the data of both sheets of different materials and different textures of real objects like chairs. We train the model for 500 epochs, where one epoch is a full pass of the data through the neural net. We randomly shuffle the training data after each epoch. Our batch size is 32 and we randomly select one batch at a time from the training data sets to feed into the neural nets. We divide the entire dataset into separate training and testing dataset to evaluate the classification network. Our MLP contains three hidden layers with each layer having 200 units and using a rectified linear unit as the activation function. The network is trained to optimize for cross-entropy loss using the Adadelta algorithm [55].

5.2 Tracking and Sensing Multiple Objects

Our discussion so far has exploited localization to isolate and find the material that objects in the environment are made of. However, sensing the materials of objects can also improve the performance of localization of these objects. To

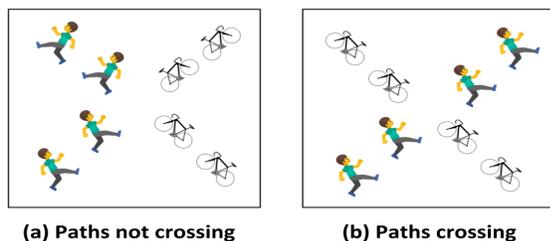


Figure 5: IntuWition can track objects with ambiguous intersecting paths by tagging objects with their respective materials.

see how, we consider a challenging problem in device-free location tracking: tracking multiple objects whose trajectories intersect.

Ambiguous Trajectories: Fig. 5 shows an example of two objects – a pedestrian and a bicycle whose trajectories intersect at a crosswalk. Traditional wireless device-free localization systems struggle to distinguish between the two trajectories in Fig. 5(a) and Fig. 5(b) – both of which result in an identical set of observed times-of-flight of reflectors.

IntuWition can resolve this challenge by exploiting a very simple property – while the location of an object can change over time, its material composition cannot. As a result, IntuWition can use the measured material composition of the two objects that are tracked to distinguish between the trajectories in Fig. 5(a) and (b). An important limitation of this approach is that it cannot disambiguate objects of precisely the same material. We believe that sensing other properties (e.g. size, shape, etc.) along with material can help address this limitation – a task for future work.

Spurious Multi-Object Reflections: A second challenge that IntuWition resolves is to eliminate spurious reflections off multiple objects before reaching the receiver. The resulting polarization observed is a combination of the properties of all materials through which the signals reflect. Specifically, the phase shift of $r_{s,s}$ and $r_{s,p}$ values of the signal received (see Sec. 3) that captures signal polarization is the sum of the $r_{s,s}$ and $r_{s,p}$ values of each material. We note that signal absorptions cause limited impact on polarization [10].

IntuWition addresses this challenge by actively modeling multi-object polarization. In particular, IntuWition sorts objects based on their observed-value of time-of-flight (i.e. a measure of distance) relative to the vehicle. It then progressively labels each object as legitimate, only marking objects as spurious if its phase shift is the sum of subsets of legitimate phase shifts. At a high level, this is because the only propagation effects that affect the horizontally and vertically polarized signals asymmetrically at the receiver are due to reflections from obstacles, and the effect of the final received

signal can be modeled as the product of the transmitted signal along with each reflector. This translates to a sum in phase effects. Thus, IntuWition quickly eliminates spurious reflections off multiple objects.

5.3 Material Properties vs. Other Effects

From atop a mobile platform, our system must deal with materials at various ranges, sizes, orientations, and textures. Our system must therefore extract the effect of material properties of the object from wireless channels from the variety of other properties of surfaces that influence channels: object size, orientation, texture and location.

Effect of Location: In handling objects of different sizes, ranges and orientations, measuring changes in polarization actually emerges as an advantageous approach when compared with measuring pure power. While reflected power could certainly distinguish a wooden and metal sheet at the same distance, it remains vulnerable to multipath and is subject to ambiguity, e.g. between a nearby wooden sheet and distant metal sheet. We instead use power ratios of the channels from a vertically polarized transmitter received by three perpendicular linearly polarized antennas. This ensures that the distance traversed to the reflector is effectively factored out of our ratio reading.

Effect of Object Size and Orientation: The challenges of varying reflector sizes and angular orientations would pose a similar issue for power-based material identification. These properties are captured in radar ranging by a property σ which they call the Radar Cross Section (RCS). Also known as the effective area, RCS is the value associated with what size an object’s reflection, independent of what its dimensions actually are, appears to a radar sensing system. Angle of orientation and physical size of this object both affect this, as intuitively, a 4’x4’ sheet oriented at 30 degrees would appear larger than one oriented at 45 degrees. Our multi-antenna approach to measuring rotations in polarization would be affected by the same effective area across all three antennas, thus removing the effect of size and angular orientation from isolating the material.

Effect of Surface Texture: We note that surface texture and material composition can be decoupled by looking at a window of reflected values, instead of a single point. Polished surfaces show a distinct peak at the time of flight of the reflected object (along with a few later peaks due to multipath effects). On the other hand, in rough volume reflectors, signals transmitted through the initial surface but have multiple opportunities to reflect through the width of the object, resulting in a shorter, wider peak.

Given the entire spectrum of possible surface textures, we could not exhaustively test this space, but to help account

for this effect, we train our machine learning models (Multi-Layer Perceptron) on diverse textures per material - such as different finishes of metal, or a wooden sheet sanded by different grit sandpaper. We present our results in 9.3.

Effect of Incident Angle: The incident angle affects polarization by the well-studied Fresnel Equations.[24]. We train our models with many different angles to help account for this effect.

6 OBJECT LOCALIZATION

Prior to sensing the material of an object-of-interest in the environment, IntuWition needs to process the received wireless channel to separate the wireless signal component reflected off this object alone. IntuWition achieves this by first finding the location of the object-of-interest, using signals from an object, collected from different perspectives, to triangulate its position. In the case of multiple objects, we remove peaks using successive interference calculation and then remove spurious multi-object reflections based on their r_{ss} and r_{sp} values, as described in 5.2. IntuWition then develops a filter to extract the wireless signal component that emerges from that particular location. The rest of this section describes IntuWition’s approach to locate surrounding objects.

6.1 Separating Signal Paths in Mobile Settings

IntuWition’s first task is to separate the signal components arriving along multiple paths as they reflect off each object in the environment. At first blush, IntuWition may achieve this using prior work on RF localization using commodity Wi-Fi radios [44]. Specifically, past systems seek to separate the time-of-flight of different signal paths from the received signal, as they reflect off various objects in the environment. An important factor that determines the resolution of measured time-of-flight of various signal paths is the available bandwidth (20 MHz for Wi-Fi). To mitigate this, past work [44] stitches together wireless channel measurements across multiple frequency bands to emulate signals from a wide-band receiver. However, an important challenge in combining frequency bands in mobile contexts such as the UAV is that the device moves in the 3-D space so that its location changes significantly between channel measurements. Hence, channels across packets change both due to change in Wi-Fi frequency as well as the change in device location between measurements.

IntuWition addresses this challenge by actively modeling both the frequency of operation and the movement of the device (e.g. on a drone) in its analysis of the wireless channels. In doing so, it retrieves both the angle-of-arrival and the time-of-flight of reflecting surfaces simultaneously. The rest of this analysis makes two simplifying assumptions for ease

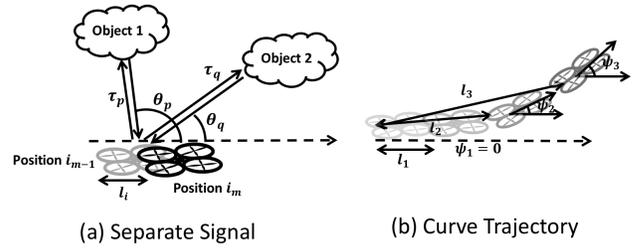


Figure 6: (a) IntuWition measures the distance and orientation of objects of interest using the wireless channels that reflect off them. (b) Notation for device (e.g. on a UAV) location along curved trajectories.

of exposition, which we relax later in this section: (1) The device moves in 2-D space; (2) phase errors due to frequency and timing offsets are accounted for.

Mathematically, assume that the device (Fig. 6) moves on a 2-D trajectory that is at a distance l_i relative to its initial position, ψ_i relative to its initial orientation and on frequency f_i when receiving packet i . We assume that the distance between the reflector and the device is much greater than the displacement of the vehicle across packets, allowing for us to approximate that the angle of the reflector does not change. To recover the time-of-flight and location of objects in the environment, IntuWition builds upon the Bartlett algorithm [26] in 2-dimensions. Specifically, we can write the power of the wireless channel received along a distance r and angle-of-arrival θ as:

$$P(r, \theta) = \left| \sum_i h_i e^{2\pi j f_i \left(\frac{2r}{c} + \frac{l_i \cos(\theta - \psi_i)}{c} \right)} \right|^2 \quad (1)$$

The above power-profile will have peaks corresponding to the polar coordinates (r, θ) of various objects in the environment relative to the vehicle. We can then extract the wireless channels corresponding to any specific object of interest at (r, θ) as:

$$h(r, \theta) = \sum_i h_i e^{2\pi j f_i \left(\frac{2r}{c} + \frac{l_i \cos(\theta - \psi_i)}{c} \right)} \quad (2)$$

We show in Sec. 5 how one can detect the material of an object at (r, θ) by analyzing $h(r, \theta)$ at the horizontally and vertically polarized antennas across frequency.

While the above analysis assumes that the available Wi-Fi frequencies and distances moved by the device between packets are uniformly spaced, in practice, this may not be true, leading to spurious peaks in $P(r, \theta)$. IntuWition builds on past work on wireless localization [44] to eliminate these peaks by leveraging the sparsity of signal multipath. Specifically, we assume that the signals that reflect off the environment emerge from a small number of dominant paths, leading to

a sparse $P(r, \theta)$ with a few distinct peaks. Mathematically:

$$\min_{\{A_p, \psi_p\}} \sum_p |A_p|, \quad (3)$$

$$\sum_i \left| h_i - \sum_p A_p w_p \right|^2 = 0 \quad (4)$$

$$w_{i,p} = e^{-2\pi j f_i \left(\tau_p + \frac{l_i \cos(\theta_p - \psi_i)}{c} \right)}$$

This optimization resembles a non-convex 2-D Non-Uniform Discrete Fourier transform (2-D NDFT). IntuWition solves it numerically using proximal gradient descent algorithms [18] generalized to 2-D to accurately localize objects, without being misled by spurious artifacts.

We make a few important generalizations of our approach: (1) While the above optimization assumes the reflector is far from the object, IntuWition solves the equivalent problem for close reflectors through a maximum likelihood approach that iterates over coordinates of closeby reflectors. As we only care about nearby objects, the search space is relatively small, enabling efficient optimization. (2) To mitigate the effect of hardware offsets, we draw from past works [40, 51] which connects two Wi-Fi chips by a wired pathway (via RF attenuators) to enable synchronization. Specifically, we connect a cable from pin 3 of the transmitter, split to pin 3 of each of the three receivers, and divide each communication channel by the reference cable channel. We then use this ratio of the wireless channels between wireless and wired pathways, across both horizontal and vertical antennas, to eliminate frequency and timing offsets. (3) It is easy to see that our analysis readily generalizes to 3-D by iterating over the polar angle ϕ_p of any reflector p in the optimization as well, i.e setting: $w_{i,p} = e^{-2\pi j f_i \left(\tau_p + \frac{l_i \cos(\theta_p - \psi_i) \sin \phi_p}{c} \right)}$.

6.2 Detecting Occluded Objects

Among the most significant challenges IntuWition faces is detecting occluded objects, whose signals are received very weakly at receiver. More problematically, the weak signals of these objects are often overwhelmed by reflections from the object in front of them. This is the classic problem of the near-far effect faced by several wireless sensing solutions [6, 34]. Past solutions address this problem in several ways for moving objects in the environment, the most common of which is background subtraction [5, 6], where these systems subtract the observed wireless channels between two time instances to filter out what changed between them.

While IntuWition can benefit from these past solutions if the device is static, they do not apply when the device moves. This is because the wireless channels between measurements change *both* due to changes in the environment and the movement of the device itself.

IntuWition addresses this challenge by developing a background subtraction algorithm that accounts for the movement of the device itself. Specifically, let us assume that $h(r, \theta)$ and $h'(r, \theta)$ are the wireless channels along any given direction (r, θ) measured at two time instances where the device is displaced by a distance Δd and re-oriented by $\Delta \theta$. It follows that any object at (r, θ) at the first time instance is now at $(r - \Delta d \cos \theta, \theta - \Delta \theta)$ (provided $r \gg \Delta d$). As a result, if the reflector at (r, θ) remain static, it is easy to see that $h(r, \theta) = h'(r - \Delta d \cos \theta, \theta - \Delta \theta)$. One can therefore perform background subtraction to detect moving object at each (r, θ) by subtracting the channels as follows:

$$\Delta h(r, \theta) = h(r, \theta) - h'(r - \Delta d \cos \theta, \theta - \Delta \theta)$$

IntuWition can then process $\Delta h(r, \theta)$ for each (r, θ) to detect the material of moving objects at each location, even if their signals are weakened due to occlusion.

6.3 Role Of Localization

Localization plays an enabling role to material sensing in our system. It is not our primary technical contribution, but localization is what allows us to isolate objects in the environment for materials classification and determines the resolution of the system. Localization also removes occlusions and determines the operational range of our system. After the object is isolated, it is passed to the material sensing algorithms described in §5 to isolate object characteristics.

7 LIMITATIONS OF INTUWITION

Given that our system is a feasibility study, we note some important limitations of IntuWition (see §9 for evaluation): It performs poorly when detecting weak reflectors whose signals are attenuated due to distance or size. Its resolution in separating multiple objects is limited by the total aggregate bandwidth of Wi-Fi: about 0.4 m in our experiments. This means that a human leaning on a wall may be misclassified. It often misses fleeting reflections due to fast moving objects. It cannot tell apart certain materials that have similar polarization characteristics or objects composed of the same material (e.g. two different humans). We are also limited in our ability to resolve a signal after significant multibounce effects, as each reflection significantly reduces the power at the receiver.

In addition, our system as a material classifier may respond most strongly to surface characteristics and misclassify clothing as humans, for example. In cases of extensive attenuation, such as around-the-corner coupled with through-wall sensing, our localization accuracy drops further. Finally, we note that because of the use of background subtraction, our system is most effective for systems where an object is introduced or where an object is moving. However, objects

of sufficient size, reflectivity, and/or visibility can still be detected without background subtraction.

8 IMPLEMENTATION

System Setup: We implement IntuWition using four WiFi cards, one for transmission and three for receiving. Each includes a linearly polarized antenna (Ettus Vert2450) connected to each Intel 5300AGN card on the Intel Galileo Boards, which run the Linux 802.11n CSI tool [16] to obtain channel state information (Fig. 7). The transmitter also has an SMA splitter and wire on the first port of the 5300AGN card that carries a reference signal to the second port of the receivers to synchronize all radios and remove phase errors caused by hardware impediments (see Sec. 6.1). The wired line is attenuated by 50 dB to avoid saturation. To measure the polarimetry in 3D space, the three receiving antennas are oriented perpendicular to each other. Note that we chose to use three Wi-Fi cards, instead of one which would suffice in principle, because we observed that the first port of the Intel 5300AGN card had a consistently higher receive gain on average when compared to the other ports (a peculiarity of the chip). The overall weight of our sensing system is 439g, which can be significantly reduced (to just the antennas) should future Wi-Fi cards on mobile platforms (e.g. onboard Wi-Fi on drones) report wireless channels as the Intel 5300 does. We place our WiFi receiver and transmitter, pictured in 7A and B, on two setups for testing: a rolling cart, with a separation of 42cm, and atop a drone as shown in 7C, with a separation of 31cm.

We note that the use of WiFi chips with a future sensing feature could reduce the chips needed to two (if one chip could accept three receivers), would eliminate the need for the Galileo boards, and could speed our channel hopping process. Further, as we are only using channel measurements for sensing, our transmissions are arbitrary packets which could instead be data for the autonomous system.

Software and Run-Time: We implement IntuWition’s algorithms in MATLAB/C++ (MLP implemented in Python) in the cloud and track location at the rate of 8 Hz and sense materials at the rate of 2 Hz. The main bottleneck in ensuring faster, real-time analysis is the amount of time it takes to sweep through all available WiFi channels in the ISM band supported by the Intel 5300 including 5-GHz bands – a problem that may be remedied by future chipsets. In its current state, a full frequency sweep for each receiver radio results in, at most, a file 18.9kB. This requires a backhaul of 37.8kBps for material sensing. For localization, the radios hop across 1/4th of the channels, so the backhaul required is the same.

Experimental Evaluation: Unless stated otherwise, our evaluation across experiments are conducted in settings as

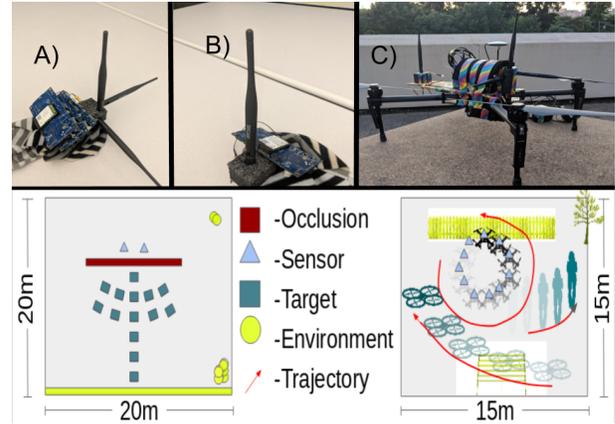


Figure 7: Setup: (A) 3 receiving and (B) 1 transmitting antennas, (C) Augmented UAV setup. Indoor (left) & outdoor testbeds (right)

described below. We evaluate the system in a 225 sq.m outdoor garage space, a 500 sq.m outdoor plaza, and four different indoor spaces ranging from 43-400 sq.m in a large university campus (see Fig. 7).

Our evaluation consists of diverse types of objects of various shapes, sizes and thicknesses. For our material tests, we use sheets of copper(3’x3’), aluminum(2’x3’, 3’x3’), birch(2’x4’), plywood(4’x2’), maple(4’x4’x.25”, 4’x4’x.5”, 4’x4’x.75”). The wooden sheets also include a variety of textures: semi-gloss, roughed up with steel wool, and sanded by 80 and 220 grit. Furthermore, we additionally test more realistic objects in our evaluation, including a chair, a table, a filing cabinet, and a car. We also recruited ten human volunteers of different body types, with half our testing performed in summer clothing (thin t-shirts and shorts) and half of it performed in a winter clothing (jackets/coats). We collected data at different orientations (40 to 150 degrees of tilt) and distances (2m to 10m away from our setup), indoors and outdoors. We perform thirty trials for each material at each configuration. All our experiments are in indoor/outdoor multipath-rich settings (walls, furniture etc.), including non line-of-sight where objects were behind wooden partitions. We emphasize that for each experiment, we train and test our system on completely different subsets of both materials and testbed locations. All of our experiments consider the presence of multiple objects whose location and material we explicitly track (up to four). We explicitly evaluate the effect of mobility of the system in Sec. 9.5, which describes a proof-of-concept evaluation of our system to sense obstacles hidden from view around a UAV. Our UAV was measured to hover at 0.05 m/s by default and moves at speeds up to 2 m/s (fast walking human). Note that graph error bars show standard deviation of location/sensing accuracy across multiple experiments.

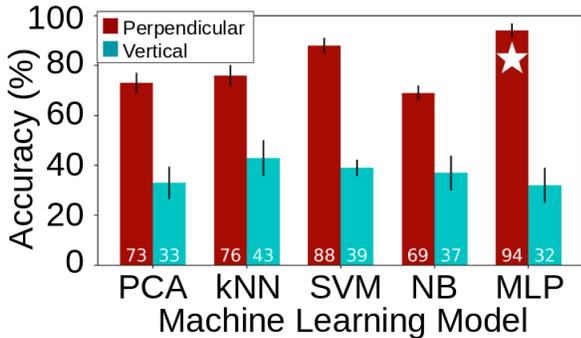


Figure 8: Three vertical antennas vs. three mutually perpendicular antennas, compared in material identification accuracy across 5 common classifiers. Note significantly higher performance for perpendicular antennas across the board, showing importance of polarization.

9 RESULTS

In this section, we describe the methods and results for our experiments. We note that localization is a precondition for material sensing, and both material sensing and localization are required for object tracking, so the results are inherently integrated but have been separated into material sensing and localization to observe trends.

9.1 Microbenchmarks

IntuWition relies on two hypotheses: that polarimetry can provide significant gain over spatial diversity as a material sensing modality, and that Multi-Layer Perceptrons are the best model for our application due to their analysis of higher-dimensional features (compared to say, kNN) and their deep architecture (allowing for more efficient and intelligent behavior than shallow ones [8].)

To test these hypotheses, we ran a preliminary experiment. To confirm polarimetry, rather than spatial diversity, was contributing to material sensing gain, we ran identical trials with three vertical receiving antennas as well as three perpendicularly oriented antennas. We collected data (specifically, power ratios of receiving antennas) in a line-of-sight lab setting to distinguish between three material types: birch, aluminum, and human, at a variety of angles and distances, to create our training set. Then, we rearranged the setup and background furniture in the lab to create the test set. Using this training and test data, we found material classification accuracy for five different common machine learning models, as shown in 8.

From this experiment, we draw two major takeaways: The use of perpendicularly oriented antennas significantly

increases material sensing accuracy across all machine learning models, and that Multi-Layer Perceptrons are indeed the ideal model for our experiment. We chose to build the Multi-Layer Perceptron to have three hidden layers based on our experimental analysis.

9.2 Object Localization

While localization is not the main focus of our paper, we include localization accuracy to isolate different materials and distances in Fig. 9 for completeness. For ground truth, we affix fiducial markers to objects and implement a camera-based fiducial tracking system using ARToolKit [2], which reports a sub-millimeter baseline accuracy. The primary takeaway from our localization results is that our system can typically isolate a line-of-sight (LOS) or non-line-of-sight (NLOS) object that is at least a meter separated from other reflectors.

9.3 Material Sensing

In this section, we discuss our core material sensing results and polarimetry’s robustness to material thickness, surface area, range, and texture.

(1) Overall Observations: In fig. 12(a), we can see IntuWition has overall classification accuracies of 95% in LOS and 92% in NLOS. The confusion matrix in fig. 11 further shows that our system tends most often to confuse objects of the same class—metals copper and aluminum, for example. We surmise our system can only distinguish them at all due to the thin layer of oxide that forms on the surface of a metal – aluminum oxide is an insulator, while copper oxide is a semiconductor, or perhaps due to polymer coatings used to prevent rusting. We don’t anticipate this confusion to be of concern in some contexts such as UAV path planning, since the exact type of metal isn’t critical to the functionality of the system.

When comparing system performance with multiple tracked objects, we see a negligible performance decrease for the two-object case, a 5% decrease in classification accuracy in the three-object case, and a 6% decrease in classification accuracy in the four-object case.

(2) Material Sensing vs. Thicknesses: We use wood to evaluate our system’s robustness to thickness, as metal only has a skin depth of roughly $1\mu\text{m}$ at Wi-Fi frequencies (mostly acting as a surface reflector) and the material composition of humans is heterogeneous.

Results: In Fig. 12(b), we observe that while 0.25" wood has less layers of water to refract compared to the other two, our system still maintains a robust classification accuracy higher than 89% even in the NLOS scenario. Further, we observe that a thickness increase of half an inch resulted in a 7% increase in classification accuracy in line-of-sight settings. We

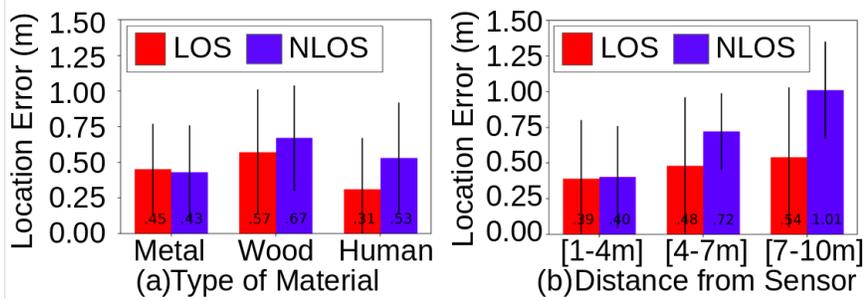


Figure 9: (a) Effect of Material Type on Localization Error: We observe our system localizes wood and metal best, which we expect is due to the larger surface areas of our material samples on average when compared to humans, for an overall mean error of 0.49m. (b) Effect of Round-trip Distance on Material Sensing: we see accuracy slightly increases with distance at the beginning, as the reflector becomes more distinguishable from the strong line-of-sight signal, and eventually falls with distance.

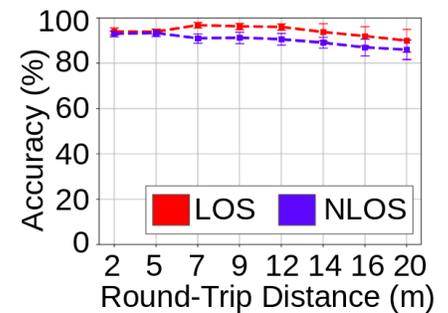


Figure 10: Effect of system to reflector one-way distance on Localization Error: we see error increase with distance, which we expect is due to less power being received from reflectors as distance increases.

True label \ Predicted label	Copper	Aluminum	Human	Plywood	Birch
Copper	90.5	9.1	0.0	0.0	0.0
Aluminum	8.5	91.5	0.0	0.0	0.0
Human	2.3	0.0	93.5	4.2	0.0
Plywood	0.0	0.0	1.9	95.0	3.1
Birch	0.0	0.0	0.0	2.9	97.1

Figure 11: This confusion matrix shows our efficacy in classifying between five different materials using our Multi-layer Perceptron Classifier.

surmise that this is because there is more available material to reflect and refract its unique polarization signature. This leads us to believe that for non-metallic objects, the material that composes the largest part of a reflector may impart the most significant share of signature, as would be necessary for identifying humans regardless of clothing.

(3) **Material Sensing vs. Surface Area:** We use a variety of material surface areas to evaluate our system’s resilience to different sizes. Specifically, we used various sizes of aluminum, copper, maple, and birch, as well as a human. These

ranged in surface area from 4ft^2 to 16ft^2 at distances ranging from 0-15 m.

Results: Fig. 12(c), shows that as surface area increases from 4ft^2 to 16ft^2 , the classification accuracy increases from 88% to 93% in NLOS and from 92% to 98% in LOS scenarios. That is, reflective surfaces with larger cross-sectional areas are easier to detect and classify, as expected. Note that lower surface areas will progressively lower accuracy (a limitation of IntuWition).

(4) **Material Sensing vs. Range:** To evaluate our system’s robustness to range, we collected measurements at up to 20 m in round-trip distance from the system to the sensed object.

Results: As shown in Fig. 10, we observe a slight increase in our system’s accuracy in the beginning, which we surmise is due to the increased ability of our system to resolve the reflector given the strong line-of-sight signal. Then, there is minimal impact in our system’s performance with up to a round-trip distance of 12 m, following which we see a gradual decrease as expected due to the reflector’s signal being received more weakly.

(5) **Material Sensing vs. Surface Texture:** Finally, to evaluate our system’s robustness to texture, we collect measurements from five surface textures of plywood of fixed size and thickness: unfinished, roughed up with steel wool, painted with semi-gloss, sanded to 80 grit (particle size of $201\ \mu\text{m}$), and sanded to 220 grit (particle size of $70\ \mu\text{m}$).

Results: As shown in Fig. 13a, the surface texture of wood does not play a statistically significant role in classification accuracy.

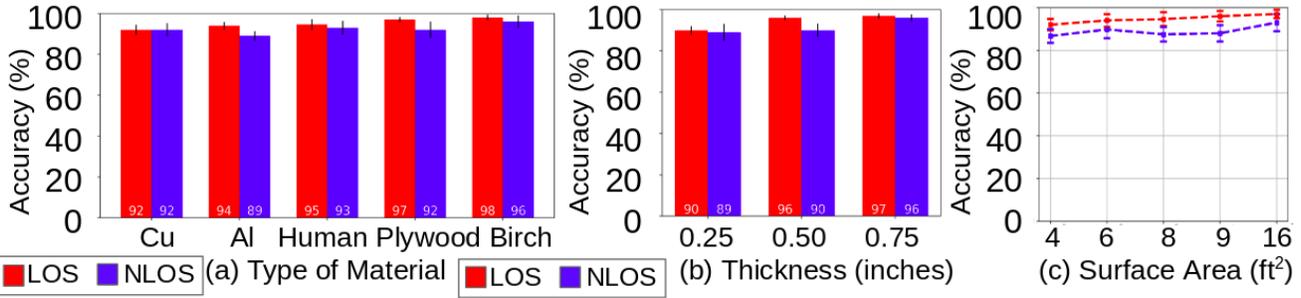


Figure 12: (a) The effect of material type on classification accuracy: we see an average accuracy of 93.5%, with higher errors for copper and aluminum (since they are more easily confused). (b) The effect of thickness of wood on classification accuracy: we see lower accuracy as thickness decreases and fewer layers of material exists to reflect its signature. (c) Material Sensing vs. Surface Area: we see accuracy increases with surface area as scattering occurs at a larger scale.

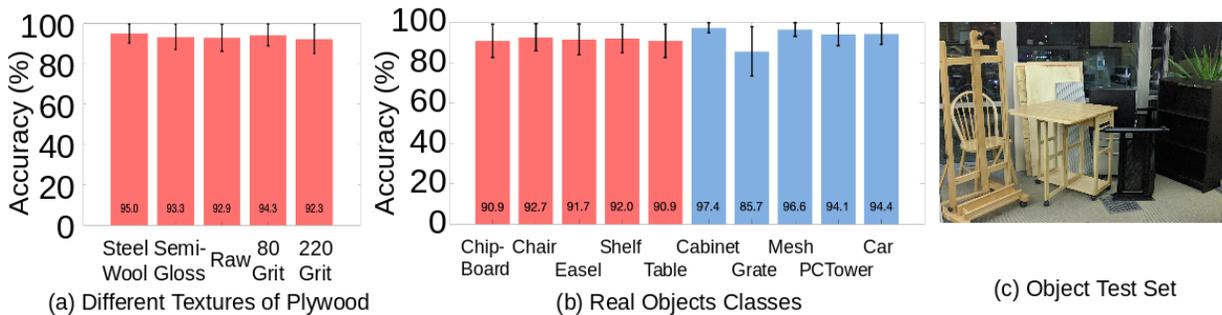


Figure 13: (a) We show our classification network accuracy dealing with different surface textures of the same material (wood). (b) We show our classification network accuracy dealing with real objects used in our daily life of different material (wood is showed as red, metal is showed as blue). (c) We show the boards, furniture, and objects used for these experiments.

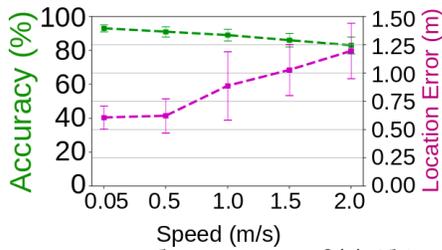


Figure 14: Measures the accuracy of (a) Object recognition (drone vs. human); and (b) Localization at varying UAV speeds.

9.4 Object Recognition of Real-Life Obstacles

To test our system’s coarse applicability to real-life objects, we train a new model to classify wood from metal using all collected test and training data up until this point to create the best model possible. We test this on new, unseen objects. These include five of each wooden and metal objects of varied material types (various wood and metal types): A rough chipboard sheet, a table, a chair, an easel, and a bookshelf; a filing cabinet, a shelving grate, a mesh furniture stand, a

matte PC tower, and the passenger-side-door of a car. These were collected in a lab setting using a 4’x4’x1” wooden sheet to occlude the setup for NLOS settings, with the exception of the car which was collected in a yet-unseen parking garage environment and was occluded by a 2’ concrete wall for NLOS readings.

Results: As seen in Fig. 13, we see relatively high accuracy across the board given that our system had only been trained on metal and wooden sheets and these new objects were unseen. We note that classification accuracy for the metals appears to be slightly higher, with the exception of the metal grate, whose structure likely introduces a lot of destructive interference.

9.5 Application: Object Recognition from a Mobile UAV

In this experiment, we demonstrate our system’s utility in avoiding two types of obstacles: humans and other UAVs for a proof-of-concept application: a delivery drone navigating indoor and outdoor spaces. We are particularly interested in capturing how the drone’s speed and vibrations affect

our system (see Sec. 5.2). Fortunately, recall that IntuWition only needs the relative trajectory of the UAV, not its absolute location at any point. Further, IntuWition applies a sliding window filter ($N=6$) on the output of inertial and motion sensors on UAVs across multiple measurements to mitigate spurious readings. Since our drone is operating around 100 revolution per second which can be translated to 100 Hz of the vibration of the drone [30], we additionally apply a high-pass filter to remove this noise. We classify between drones and humans from atop a drone at a few different speeds: 0, 0.5, 1.0, 1.5 and 2.0 meter per second along a circular trajectory with a 5 m radius. Note that we do not test at higher speeds owing to limitations of the UAV platform and to ensure user safety. We further localize the reflector using 3-D MUSIC to measure localization error in three dimensions.

Results: We notice that increasing UAV speed decreases the accuracy of material sensing. One hypothesis is that because the rate of sweeping Wi-Fi frequencies is around 2 Hz, increased vibration of the propellers impacts our error in localization which in turn influences material sensing accuracy. Our MLP model has a prediction latency around 0.04 second (25 Hz). We observe that our accuracy falls off with speed, but we are still able to achieve 89% mean accuracy in classifying between humans and drones at speeds of up to 2 m/s and mean localization accuracy of 0.87 m.

10 CONCLUSION AND FUTURE WORK

This paper presents IntuWition, a system that explores sensing the material and location of hidden objects in the environment using commodity Wi-Fi radios. IntuWition analyzes the change in polarization of the wireless signals as they reflect off different objects surrounding a Wi-Fi device to infer their material composition. A detailed evaluation demonstrates promising accuracy in both localization and material identification.

While IntuWition categorizes between five types of materials, we believe future work can take this much further. We believe high-bandwidth radios such as 802.11ad can greatly improve material sensing resolution. We also leave for future work developing optimized hardware that performs all material processing on-board the device, as opposed to the cloud. In the context of UAVs, fusing information with on-board cameras and dealing with object recognition system design challenges on robotic platforms remains an important task for future work.

11 ACKNOWLEDGEMENTS

This work has been supported in part by the National Science Foundation through CRII Grant #1657318, as well as the ARCS Foundation, the Michel E. and Kathy Doreau Fellowship in Electrical and Computer Engineering and the

Marshall Fellowship. We further thank our labmates at the WiTech lab for their continued support, and the many interns who have worked on this project, including Namrita Murali, Isha Iyer, Atul Bansal, Zijian Chen, and Raghav Sonavane.

REFERENCES

- [1] [n. d.]. Airborne Detect And Avoid radar for small to medium UAVs enables safe beyond visual line of sight operations. <https://echodyne.com/wp-content/uploads/2017/05/MESA-DAA-Product-Sheet.pdf>. (Accessed on 08/05/2018).
- [2] [n. d.]. ARToolKit v5.x. (Accessed on 08/06/2018).
- [3] [n. d.]. MESA-X-EVU Evaluation Unit Datasheet. http://xponential.vporoom.com/Echodyne/download/MESA-X_EVU_2016_0428.pdf
- [4] [n. d.]. Military Radar Technology. <https://sputniknews.com/news/201504071020549211/>. (Accessed on 06/23/2018).
- [5] Fadel Adib, Zachary Kabelac, Dina Katabi, and Robert C Miller. 2014. 3D Tracking via Body Radio Reflections.. In *NSDI*, Vol. 14. 317–329.
- [6] Fadel Adib and Dina Katabi. 2013. *See through walls with WiFi!* Vol. 43. ACM.
- [7] Kamran Ali, Alex X Liu, Wei Wang, and Muhammad Shahzad. 2015. Keystroke recognition using wifi signals. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 90–102.
- [8] Yoshua Bengio, Yann LeCun, et al. [n. d.]. Scaling learning algorithms towards AI. ([n. d.]).
- [9] Xiaolin Bian, Yun Shao, Wei Tian, Shiang Wang, Chunyan Zhang, Xiaochen Wang, and Zhixin Zhang. 2017. Underwater Topography Detection in Coastal Areas Using Fully Polarimetric SAR Data. *Remote Sensing* 9, 6 (2017).
- [10] Chi-Chih Chen, M. B. Higgins, K. O’Neill, and R. Detsch. 2001. Ultrawide-bandwidth fully-polarimetric ground penetrating radar classification of subsurface unexploded ordnance. *IEEE Transactions on Geoscience and Remote Sensing* 39, 6 (June 2001), 1221–1230. <https://doi.org/10.1109/36.927444>
- [11] S. Depatla and Y. Mostofi. 2018. Crowd Counting Through Walls Using WiFi. In *2018 IEEE International Conference on Pervasive Computing and Communications (PerCom)(PERCOM)*, Vol. 00. 1–10. <https://doi.org/10.1109/PERCOM.2018.8444589>
- [12] S. Depatla, A. Muralidharan, and Y. Mostofi. 2015. Occupancy Estimation Using Only WiFi Power Measurements. *IEEE Journal on Selected Areas in Communications* 33, 7 (July 2015), 1381–1393. <https://doi.org/10.1109/JSAC.2015.2430272>
- [13] Ashutosh Dhekne, Mahanth Gowda, Yixuan Zhao, Haitham Hassanieh, and Romit Roy Choudhury. 2018. LiquiID: A Wireless Liquid Identifier. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys ’18)*. ACM, New York, NY, USA, 442–454. <https://doi.org/10.1145/3210240.3210345>
- [14] Mario Gerla, Eun-Kyu Lee, Giovanni Pau, and Uichin Lee. 2014. Internet of vehicles: From intelligent grid to autonomous cars and vehicular clouds. In *Internet of Things (WF-IoT), 2014 IEEE World Forum on*. IEEE, 241–246.
- [15] Daniel Göhring, Miao Wang, Michael Schnürmacher, and Tinosch Ganjineh. 2011. Radar/lidar sensor fusion for car-following on highways. In *Automation, Robotics and Applications (ICARA), 2011 5th International Conference on*. IEEE, 407–412.
- [16] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. 2011. Tool release: Gathering 802.11 n traces with channel state information. *ACM SIGCOMM Computer Communication Review* 41, 1 (2011), 53–53.
- [17] Jiaming Hao, Yu Yuan, Lixin Ran, Tao Jiang, Jin Au Kong, C. T. Chan, and Lei Zhou. 2007. Manipulating Electromagnetic Wave Polarizations

- by Anisotropic Metamaterials. *Phys. Rev. Lett.* 99 (Aug 2007), 063908. Issue 6. <https://doi.org/10.1103/PhysRevLett.99.063908>
- [18] Ke Hou, Zirui Zhou, Anthony Man-Cho So, and Zhi-Quan Luo. 2013. On the linear convergence of the proximal gradient method for trace norm regularization. In *Advances in Neural Information Processing Systems*. 710–718.
- [19] Albert S Huang, Abraham Bachrach, Peter Henry, Michael Krainin, Daniel Maturana, Dieter Fox, and Nicholas Roy. 2017. Visual odometry and mapping for autonomous flight using an RGB-D camera. In *Robotics Research*. Springer, 235–252.
- [20] Donny Huang, Rajalakshmi Nandakumar, and Shyamnath Gollakota. 2014. Feasibility and limits of wi-fi imaging. *Proceedings of the ACM SenSys* (2014), 266–279. <https://doi.org/10.1145/2668332.2668344>
- [21] Ebi Jose and Martin D Adams. [n. d.]. An Augmented State SLAM formulation for Multiple Line-of-Sight Features with Millimetre Wave RADAR. ([n. d.]). <https://pdfs.semanticscholar.org/2dc4/28d47ce57715858aa2c288557b97a4dcd034.pdf>
- [22] C. R. Karanam, B. Korany, and Y. Mostofi. 2018. Magnitude-Based Angle-of-Arrival Estimation, Localization, and Target Tracking. In *Proceedings of ACM/IEEE International Conference on Information Processing in Sensor Networks*. 254–265.
- [23] Chitra R Karanam and Yasamin Mostofi. 2017. 3D through-wall imaging with unmanned aerial vehicles using WiFi. In *Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks*. ACM, 131–142.
- [24] D. V. Kolyadov, L. P. Ligthart, and A. I. Kozlov. 2006. Amplitude-phase method allowing the determination of the complex dielectric permittivity of underlying surfaces using polarimetric radar remote sensing. In *2006 European Radar Conference*. 142–145. <https://doi.org/10.1109/EURAD.2006.280294>
- [25] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. *SIGCOMM Comput. Commun. Rev.* 45, 4 (Aug. 2015), 269–282. <https://doi.org/10.1145/2829988.2787487>
- [26] Swarun Kumar, Stephanie Gil, Dina Katabi, and Daniela Rus. 2014. Accurate indoor localization with zero start-up cost. *Proceedings of the 20th annual international conference on Mobile computing and networking - MobiCom '14* (2014), 483–494. <https://doi.org/10.1145/2639108.2639142>
- [27] Swarun Kumar, Ezzeldin Hamed, Dina Katabi, and Li Erran Li. 2014. LTE radio analytics made easy and accessible. In *ACM SIGCOMM Computer Communication Review*, Vol. 44. ACM, 211–222.
- [28] Swarun Kumar, Lixin Shi, Nabeel Ahmed, Stephanie Gil, Dina Katabi, and Daniela Rus. 2012. Carspeak: a content-centric network for autonomous driving. *ACM SIGCOMM Computer Communication Review* 42, 4 (2012), 259–270.
- [29] W. S. L. Lee, C. Fumeaux, and W. Withayachumnankul. 2018. Metasurfaces for Terahertz Polarimetry. In *2018 IEEE Asia-Pacific Conference on Antennas and Propagation (APCAP)*. 514–516. <https://doi.org/10.1109/APCAP.2018.8538085>
- [30] Zhenming Li, Mingjie Lao, Swee King Phang, Mohamed Redhwan Abdul Hamid, Kok Zuea Tang, and Feng Lin. [n. d.]. Development and Design Methodology of an Anti-Vibration System on Micro-UAVs. ([n. d.]).
- [31] Luis Mejias, Scott McNamara, John Lai, and Jason Ford. 2010. Vision-based detection and tracking of aerial targets for UAV collision avoidance. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*. IEEE, 87–92.
- [32] Phuc Nguyen, Hoang Truong, Mahesh Ravindranathan, Anh Nguyen, Richard Han, and Tam Vu. 2017. Matthan: Drone Presence Detection by Identifying Physical Signatures in the Drone’s RF Communication. In *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '17)*. ACM, New York, NY, USA, 211–224. <https://doi.org/10.1145/3081333.3081354>
- [33] Jaan Praks, Elise Colin Koeniguer, and Martti T. Hallikainen. 2009. Alternatives to target entropy and alpha angle in SAR polarimetry. *IEEE Transactions on Geoscience and Remote Sensing* 47, 7 (2009), 2262–2274. <https://doi.org/10.1109/TGRS.2009.2013459>
- [34] Qifan Pu, Sidhant Gupta, Shyamnath Gollakota, and Shwetak Patel. 2013. Whole-home gesture recognition using wireless signals. In *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 27–38.
- [35] Hang Qiu, Fawad Ahmad, Fan Bai, Marco Gruteser, and Ramesh Govindan. 2018. AVR: Augmented Vehicular Reality. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '18)*. ACM, New York, NY, USA, 81–95. <https://doi.org/10.1145/3210240.3210319>
- [36] Subramanian Ramasamy, Roberto Sabatini, Alessandro Gardi, and Jing Liu. 2016. LIDAR obstacle warning and avoidance system for unmanned aerial vehicle sense-and-avoid. *Aerospace Science and Technology* 55 (2016), 344–358.
- [37] C. R. Karanam S. Depatla and Y. Mostofi. 2017. Robotic Through-Wall Imaging. *IEEE Antenna and Propagation Magazine, special issue on Electromagnetic Inverse Problems for Sensing and Imaging* (April 2017).
- [38] D. L. Schuler, J. S. Lee, T. L. Ainsworth, and M. R. Grunes. 2000. Terrain topography measurement using multipass polarimetric synthetic aperture radar data. *Radio Science* 35, 3 (May 2000), 813–832. <https://doi.org/10.1029/1999RS002266>
- [39] D. L. Schuler, Jong-Sen Lee, and G. De Grandi. 1996. Measurement of topography using polarimetric SAR images. *IEEE Transactions on Geoscience and Remote Sensing* 34, 5 (Sep 1996), 1266–1277. <https://doi.org/10.1109/36.536542>
- [40] Souvik Sen, Romit Roy Choudhury, and Srihari Nelakuditi. 2010. CSMA/CN: Carrier Sense Multiple Access with Collision Notification. In *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking (MobiCom '10)*. ACM, New York, NY, USA, 25–36. <https://doi.org/10.1145/1859995.1859999>
- [41] Longfei Shangguan and Kyle Jamieson. 2016. Leveraging Electromagnetic Polarization in a Two-Antenna Whiteboard in the Air. *Proceedings of the 12th International Conference on Emerging Networking Experiments and Technologies* (2016), 443–456. <https://doi.org/10.1145/2999572.2999601>
- [42] Jiancheng Shi, James Wang, Ann Y Hsu, Peggy E O’Neill, and Edwin T Engman. 1997. Estimation of bare surface soil moisture and surface roughness parameter using L-band SAR image data. *IEEE Transactions on Geoscience and Remote Sensing* 35, 5 (1997), 1254–1266.
- [43] Suraj Srinivas, Ravi Kiran Sarvadevabhatla, Konda Reddy Mopuri, Nikita Prabhu, Srinivas SS Kruthiventi, and R Venkatesh Babu. 2016. A taxonomy of deep convolutional neural nets for computer vision. *Frontiers in Robotics and AI* 2 (2016), 36.
- [44] Deepak Vasisht, Swarun Kumar, and Dina Katabi. 2016. Decimeter-Level Localization with a Single WiFi Access Point. *NsdI 2016 ii* (2016), 165–178. <https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/vasisht>
- [45] Jue Wang and Dina Katabi. 2013. Dude, where’s my card?: RFID positioning that works with multipath and non-line of sight. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 51–62.
- [46] 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 (MobiCom '17)*. ACM, New York, NY, USA, 288–300. <https://doi.org/10.1145/3117811.3117830>

- [47] Yan Wang, Jian Liu, Yingying Chen, Marco Gruteser, Jie Yang, and Hongbo Liu. 2014. E-eyes: Device-free Location-oriented Activity Identification Using Fine-grained WiFi Signatures. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking (MobiCom '14)*. ACM, New York, NY, USA, 617–628. <https://doi.org/10.1145/2639108.2639143>
- [48] Kaishun Wu. 2016. Wi-metal: Detecting metal by using wireless networks. In *Communications (ICC), 2016 IEEE International Conference on*. IEEE, 1–6.
- [49] Yaxiong Xie, Zhenjiang Li, and Mo Li. 2015. Precise Power Delay Profiling with Commodity WiFi. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom '15)*. ACM, New York, NY, USA, 53–64. <https://doi.org/10.1145/2789168.2790124>
- [50] Yaxiong Xie, Jie Xiong, Mo Li, and Kyle Jamieson. 2016. xD-track: Leveraging Multi-dimensional Information for Passive Wi-fi Tracking. In *Proceedings of the 3rd Workshop on Hot Topics in Wireless (HotWireless '16)*. ACM, New York, NY, USA, 39–43. <https://doi.org/10.1145/2980115.2980127>
- [51] Jie Xiong, Karthikeyan Sundaresan, and Kyle Jamieson. 2015. Tone-track: Leveraging frequency-agile radios for time-based indoor wireless localization. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 537–549.
- [52] Zhicheng Yang, Parth H Pathak, Yunze Zeng, Xixi Liran, and Prasant Mohapatra. 2016. Monitoring vital signs using millimeter wave.. In *MobiHoc*. 211–220.
- [53] 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*. ACM, 833–841.
- [54] A. Zakharov, L. Zakharova, M. Sorochinsky, and T. Chimitdorzhiev. 2018. SAR Polarimetry in Remote Sensing of Arctic Region. In *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*. 2394–2397. <https://doi.org/10.1109/IGARSS.2018.8518085>
- [55] Matthew D. Zeiler. 2012. ADADELTA: An Adaptive Learning Rate Method. *CoRR* abs/1212.5701 (2012). arXiv:1212.5701 <http://arxiv.org/abs/1212.5701>
- [56] Yanzi Zhu, Yibo Zhu, Ben Y. Zhao, and Haitao Zheng. 2015. Reusing 60GHz Radios for Mobile Radar Imaging. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom '15)*. ACM, New York, NY, USA, 103–116. <https://doi.org/10.1145/2789168.2790112>