

SiWa: See into Walls via Deep UWB Radar

Tianyue Zheng^{1*}, Zhe Chen^{1*}, Jun Luo¹, Lin Ke², Chaoyang Zhao³, Yaowen Yang³

¹School of Computer Science and Engineering, Nanyang Technological University, Singapore

²Institute of Materials Research and Engineering, A*STAR, Singapore

³School of Civil and Environmental Engineering, Nanyang Technological University, Singapore

Email: {tianyue002, junluo}@ntu.edu.sg, chenz@ssijri.com, karen-kl@imre.a-star.edu.sg, {cy.zhao, cywyang}@ntu.edu.sg

ABSTRACT

Being able to *see into walls* is crucial for diagnostics of building health; it enables inspections of wall structure without undermining the structural integrity. However, existing sensing devices do not seem to offer a full capability in mapping the in-wall structure while identifying their status (e.g., seepage and corrosion). In this paper, we design and implement SiWa as a low-cost and portable system for wall inspections. Built upon a customized IR-UWB radar, SiWa scans a wall as a user swipes its probe along the wall surface; it then analyzes the reflected signals to synthesize an image and also to identify the material status. Although conventional schemes exist to handle these problems individually, they require troublesome calibrations that largely prevent them from practical adoptions. To this end, we equip SiWa with a deep learning pipeline to parse the rich sensory data. With innovative construction and training, the deep learning modules perform structural imaging and the subsequent analysis on material status, without the need for repetitive parameter tuning and calibrations. We build SiWa as a prototype and evaluate its performance via extensive experiments and field studies; results evidently confirm that SiWa accurately maps in-wall structures, identifies their materials, and detects possible defects, suggesting a promising solution for diagnosing building health with minimal effort and cost.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing systems and tools.**

KEYWORDS

In-wall imaging, material identification, commercial-grade RF-sensing, IR-UWB, Synthetic Aperture Radar (SAR).

ACM Reference Format:

T. Zheng, Z. Chen, J. Luo, L. Ke, C. Zhao, and Y. Yang. 2022. SiWa: See into Walls via Deep UWB Radar. In *The 27th Annual International Conference on Mobile Computing and Networking (ACM MobiCom'21)*, January 31-February 4, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3447993.3483258>

* Both authors contributed equally to this research.

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.

ACM MobiCom'21, January 31-February 4, 2022, New Orleans, LA, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-8342-4/22/01...\$15.00

<https://doi.org/10.1145/3447993.3483258>

1 INTRODUCTION

Walls in modern architectures are becoming increasingly complex due to functional [27, 74], structural [47, 87], and aesthetic [33, 77] advancements. However, the intricate and complex structures inside these walls may become deteriorated and even damaged as buildings age. Whereas breaking and reconstructing a damaged wall (or part) is certainly the last resort, a better maintenance method is to detect in-wall conditions non-destructively [30, 31, 51]. Presumably, it would be most helpful to open a “virtual window” into walls to expose the structure inside. Aided by such “visual” cues, we may diagnose the healthiness of the wall structure without unnecessary destruction that can be dangerous and costly. In addition, we may detect leakage pathways and rebar (reinforcing bar) corrosion at their early stage, hence repairing them with a less destructive method [60, 68].

One promising approach to non-destructive structural inspection is *radio frequency (RF)-sensing*. In fact, the wireless sensing community has witnessed a rapid growth of sensing capability enabled by RF signals in the past decade. By far, several systems have been implemented to “see” the location, shape, and motion of objects [2, 9, 24, 36, 38, 50, 81, 88, 91], as well as their material compositions [50, 85]. By leveraging the penetration and diffraction of RF signals, some of these proposals deliver an “X-ray”-like capability to see *through* certain types of occlusions (e.g., drywall or fog) [2, 3, 24, 38, 81]. However, they lack specific capabilities of seeing *into* walls (especially those built by concrete and bricks), which require processing schemes distinct from those commonly applied to RF signals propagating mainly through air.

Designing a fully functional system to see into walls is very challenging. First of all, low-cost yet high-performance hardware simply does not exist. While cheap entry-level Stud Finders [64] provide only rudimentary functionalities, professional Ground-Penetrating Radars (GPRs) [70, 73] deliver imaging results at a cost over 10kUSD, yet they penetrate only asphalt and offer no material identification ability. Second, though new hardware (e.g., the IR-UWB radar adopted by us [55]) have emerged in recent years, leveraging signals propagating in walls for diagnosis purpose is highly non-trivial, as RF-related research has focused mainly on air-propagating signals [16, 24, 50, 71, 84, 85, 90, 91], leaving us with only a handful of experimental algorithms [7, 18, 62, 66]. Last but not least, experimental algorithms often fail when facing realistic situations. For example, excessive reflection at the rough air-wall boundary can submerge useful signals, while the heterogeneity of walls (e.g., material and depth) may require algorithms to be substantially calibrated for every new situation.

To tackle these challenges, we propose **SiWa** for **Seeing into Walls**. Constructed based on a commodity-grade IR-UWB radar [55],



This is the author's version of the work. It is posted at <https://tianyue.github.io/pubs/siwa.pdf> for your personal use. Not for redistribution. The definitive version was published in *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, pp. 323–336, 2021, doi: 10.1145/3447993.3483258.




Figure 1: SiWa collects RF reflections as a user swipes its probe across a wall surface, then it intelligently processes the signals for in-wall structural imaging and material identification (e.g. rebar layout and if a rebar is corroded or not).

SiWa can be held in hand to swipe across a wall surface. It processes the reflected signals by emulating *synthetic aperture radar* (SAR) and *spectrometer* for in-wall structural imaging and material identification, as illustrated in Figure 1. In particular, SiWa aims to i) detect in-wall structures (e.g., rebar) and defects (e.g., cracks) via imaging, and ii) recognize their status (e.g., whether a rebar is corroded or a crack contains water seepage) via material identification. In order to avoid the complicated calibration of model-based approaches, SiWa exploits an end-to-end deep learning pipeline to process signals in a self-adaptive manner. Our major contributions in designing and implementing SiWa are summarized as follows:

- We design a low-cost, easy-to-carry prototype with a competitive *see into wall* performance. Leveraging an IR-UWB radar to achieve both cost-effectiveness and compactness, we exploit the radar’s wide bandwidth and extend it with dual-mode polarization antennas to obtain high imaging resolution and classification accuracy.
- We conduct detailed investigations on existing model-based techniques for RF-based structural imaging and material identification. Whereas these two functionalities have always been treated separately, our study indicates that combining them into a single RF-sensing device is possible, but model-based approaches fail to handle realistic sensing tasks.
- We equip SiWa with a deep learning pipeline to extract useful information from the rich sensory data delivered by SiWa hardware. This pipeline is designed to perform in-wall structural imaging and material identification in a self-adaptive manner. Essentially, SiWa exploits novel encoder-decoder and adversarial learning to fulfill its tasks robustly when facing diverse environments, which, to the best of our knowledge, have never been achieved in a non-destructive way.
- We thoroughly evaluate SiWa with extensive experiments and field studies. The promising results show that SiWa can indeed see into walls, so as to generate high-quality structural imaging, identify materials, and diagnose potential problems with high accuracy.

The overall idea of SiWa is to showcase the potential of RF-sensing in realizing an important application. Therefore, SiWa (albeit being a fully functional prototype) is by no means close to a final product, as its many design aspects are still tentative and experimental. The rest of the paper is organized as follows. Section 2 motivates

our design by introducing the background for in-wall structural imaging and material identification. Section 3 presents the system design of SiWa, with detailed implementation discussed in Section 4. Section 5 reports the evaluation results. Technical limitations are discussed in Section 6. Finally, Section 7 concludes this paper and points out future directions.

2 BACKGROUND AND MOTIVATIONS

In this section, we provide a brief background on in-wall imaging and material identification. Detection of subsurface targets is first studied by geologists [11], physicists [25], and chemists [61]. They mainly employ different physical and (electro-)chemical properties for the detection of targets. However, these methods only partially solve the problem of seeing into walls, so we hereafter focus on schemes that do have certain practical adoptions.

2.1 Destructive and Intrusive Methods

In order to know the structures inside a specimen, destructive



Figure 2: Drilling for sensor embedding.

methods are intuitive, and the results of them are easier to interpret [26, 86]. However, since these methods require the specimen to be destroyed, they are only suitable for objects that are mass-produced and not currently in use. For large building structures still in use, inspection by destruction is clearly not a sensible approach. As a more reasonable method, intrusive test-

ing encompasses embedding sensors in a wall [17, 19] and sticking electrodes on a wall [34, 76]. However, embedding sensors into the wall and applying electrodes with conductive fluids can be costly and also degrade the structural integrity of the wall [12].

2.2 Wireless Sensing Methods

To avoid the destructive and intrusive nature of the aforementioned methods, people have been seeking alternative methods to see into walls for a long time. Among them, various wireless methods employing acoustic and electromagnetic waves are most worthy of investigation. In the following, we discuss the pros and cons of these methods and also perform a comparative study.

2.2.1 Acoustic Sensing. Although proposals on non-destructive acoustic imaging [42, 52] and a small number of industrial products are also available [1], this technology is not suitable for our purpose because of the following reasons. First of all, acoustic signals are easily corrupted by interference, thus inapt for construction sites commonly filled with interference. Second, generating acoustic waves that propagate transversely requires sophisticated antenna arrays, rendering it very hard to support polarimetry for material identification as described in [85]. Third, most commercial acoustic transceivers are narrowband, hence the reflected signals barely carry any dispersion information pertinent to characterizing materials. Last but not least, generating wideband acoustic signals with commercial transceivers can be impractical since impedance mismatch [45] caused by a high frequency largely prevents the waves to be delivered into concrete walls.

2.2.2 Ground Penetrating Radar (GPR). As discussed in Section 1, professional GPR systems with imaging capabilities are rather expensive and mainly used for penetrating asphalt pavement [70]. Moreover, these devices are closed source and provide no Application Programming Interface (API), so it is nearly impossible to extend them to achieve material identification and structural diagnosis, which are indispensable capabilities for actually inspecting the status of in-wall structures.



Figure 3: GPR.

2.2.3 Commodity-grade RF Sensing. Commodity-grade RF devices strike a good balance between cost and performance. Some examples include Wi-Fi, IR-UWB radar, and mmWave radar operating at different frequencies. To select an appropriate device in serving our purpose of seeing into walls, we perform a preliminary study on the capabilities of these devices. We embed a rebar into a concrete wall at a depth of 5cm and swipe these devices over the wall surface. The received signals are visualized in Figure 4. We observe that, though Wi-Fi successfully penetrates the wall, excessive diffraction makes it unable to focus on and locate the rebar. Moreover, the narrow bandwidth of 20MHz confines the spatial resolution of Wi-Fi. IR-UWB radar successfully penetrates the wall and forms a typical “hyperbola” SAR pattern of a reflector. Two FMCW radars at 24GHz and 77GHz cannot penetrate into the wall; the transmitted signals are mostly reflected back by the air-wall boundary. We further test non-contact scanning performance of the IR-UWB radar. Figure 5a and 5b show the received signals when the probe is swiped respectively at 5 cm and 50 cm from the wall. Compared with the result in Figure 4b, reflection of the wall becomes much stronger and the hyperbola becomes less evident at a shorter distance, but totally disappears at a further distance.

Here are a few takeaways from this preliminary study. First, bistatic devices (e.g., Wi-Fi) whose transmitter (tx) gets separated from receiver (rx) are not suitable for our purpose anyway. On

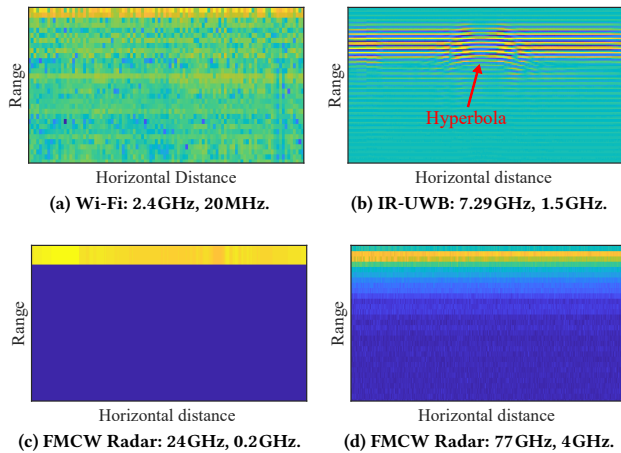


Figure 4: Preliminary study of four RF devices, with their center frequencies and bandwidths specified.




Figure 5: Non-contact scanning is vulnerable to excessive loss at the air-wall boundary and ambient interference.

one hand, the results obtained by putting tx and rx Wi-Fi antennas on the same side of the wall fail miserably compared even with Figure 4a (which is obtained by leaving the two antennas on different sides). On the other hand, walls with single-side access (e.g., in a tunnel) demand monostatic devices with co-located tx and rx antennas. Second, there is a trade-off in selecting a suitable center frequency: a low frequency helps penetration but may also cause too excessive diffraction to miss targets, while a high frequency can substantially limit the penetration ability. Third, a wider bandwidth is preferred to achieve a higher spatial resolution and imaging quality. Finally, a contact design swiping across wall surface is preferred over a non-contact approach scanning from a distance, because the latter faces excessive loss at the air-wall boundary and ambient interference. Considering the above factors, it is evident that the IR-UWB radar of 7.29GHz center frequency and 1.5GHz bandwidth has the best overall performance, when used in a contact manner.

Remarks. Because IR-UWB and FMCW are equivalent as a time-frequency dual pair [16], we choose IR-UWB to be the default technology in the following discussions. We do not consider static radar imaging [3, 78] because SAR offers flexible extension in detection area, higher SNR after coherent integration within a synthetic aperture [8], and lower cost brought by its compact design [75].

3 SIWA DESIGN

This section introduces the design of SiWa, whose rough diagram is shown in Figure 6. The hardware front-end (including an IR-UWB chip and antennas) collects SAR data matrix containing two-channel polarized and wideband signals. The SAR data matrix is then leveraged by the back-end software for in-wall imaging, which subsequently helps to locate targets and in turn identifies their respective material and status.




Figure 6: Overview of SiWa's design.

3.1 Modeling Radar Signal

By scanning across a wall surface, SiWa works as a SAR that reconstructs a 2D image of the in-wall structure by synthesizing the reflected signals gathered at multiple positions. To model the process, we first have a look at the UWB baseband signal taking the




Figure 7: Formation of the 2-D signal matrix r .

form of a Gaussian pulse:

$$s(t) = \alpha_{\text{tx}} \exp\left(-t^2/2\sigma^2\right), \quad (1)$$

and this baseband signal is then modulated onto a carrier to get:

$$x(t) = s(t) \cdot \cos(2\pi f_c t), \quad (2)$$

where α_{tx} is the pulse amplitude, σ^2 is the variance defined by the -10 dB bandwidth, and f_c is the carrier frequency,

Assuming K (point scatterer) targets present in a wall, the tx signal travels to the k -th target and its echo scatters back after a time delay $\tau_k = \frac{2r_k}{c}$, where r_k is the range to the k -th target, and c is light speed in the wall. So the rx signal becomes:

$$r(t) = \sum_{k=1}^K \beta_k s(t - \tau_k) \cdot \cos(2\pi f_c(t - \tau_k)) + n(t), \quad (3)$$

where β_k is a scaling term attributed to propagation and reflection, and $n(t)$ represents the Gaussian noise. A radar repeats this tx-rx process at a regular interval to generate 1-D sequences shown in Eq. (3). By stacking these 1-D signals one by one, we get a 2-D signal matrix shown in Figure 7. Generalizing Eq. (3), we use $r(x, t)_{z=0}$ to represent the 2-D signal matrix, where x denotes the moving distance of the radar, t represents the *range dimension*, and $z = 0$ indicates that the radar stays at zero depth. As we slide the radar on the wall surface, this signal matrix captures information for reconstructing images of in-wall structures. This model is also valid for extended scatterers, because the received signal can be represented by replacing the summation in Eq. (3) with an integration for a scatterer extending within the $x-z$ plane; otherwise, as the radar antenna has a fixed beamwidth, a scatterer extending out of the $x-z$ plane is illuminated for a finite “thickness” and can be approximated as a point scatterer in the imaging plane [58, 72].

3.2 A Conventional Design

We first study how conventional signal processing methods leverage $r(x, t)_{z=0}$ to realize functions required by SiWa. It serves as an inspiration for our deep learning-based design elaborated later.

3.2.1 Imaging by SAR. As we move the radar, the distance from a target located at (x_0, z_0) first decreases and then increases. On the received 2-D signal matrix, the trace of the Gaussian pulse forms a hyperbola as shown in Figure 7, which can be located along range dimension by $t(x, x_0, z_0) = \frac{2}{c} \sqrt{(x - x_0)^2 + z_0^2}$ [7]; an example of the trace is shown in Figure 4b. However, raw data in this form are often hard to interpret since a point target emerges as a hyperbola, and multiple targets may get mixed up. Existing proposals translating (temporal) range t to (geometric) depth z include the frequency-domain Range Migration Algorithm (RMA) [7, 66] and

time-domain back-projection algorithm [15, 54]; they all aim to “refocus” the energy spread on the hyperbola. Among them, the time-wavenumber domain implementation of RMA becomes the de-facto standard, thanks to its higher computing efficiency achieved by Fast Fourier Transform (FFT); we briefly explain the workflow and weaknesses of RMA in the following.

A matched filter is firstly applied along the range dimension of the 2-D signal matrix to achieve range compression, i.e., reducing the width of the transmitted Gaussian pulse, so as to enhance range resolution. Once filtering is completed, a 2-D FFT is performed to transform the matrix into the frequency-wavenumber domain:

$$R(\kappa_x, \omega)_{z=0} = \iint r(x, t)_{z=0} \cdot e^{-j(\omega t + \kappa_x x)} dx dt, \quad (4)$$

where ω is the frequency conjugate to the arrival time t , and analogously, κ_x is *wavenumber* (i.e., the spatial frequency) conjugate to the *moving dimension* x . κ_x is related to ω and κ_z (wavenumber in depth dimension z) by $\omega = \frac{c}{2} \cdot \sqrt{\kappa_x^2 + \kappa_z^2}$. In order to focus the energy in the current domain, we trace back the wave propagation in depth z to obtain $e^{j\kappa_z z} \cdot R(\kappa_x, \omega)_{z=0}$. As a result, the focused signal at an arbitrary depth z in moving-depth domain can be recovered by a time reversal of $t_0 = 2z_0/c$ and an inverse FFT:

$$\tilde{r}(x, z)_{t=-t_0} = \frac{1}{(2\pi)^2} \iint R(\kappa_x, \omega)_{z=0} \cdot e^{j(\kappa_z z - \omega t_0)} e^{j\kappa_x x} d\omega d\kappa_x,$$

and the resulting $\tilde{r}(x, z)_{t=-t_0}$ is indeed the image on $x-z$ plane.

This seemingly straightforward process of RMA actually requires several parameters to be specified in advance. For example, the speed v of the moving probe determines the eccentricity of the hyperbola, and the permittivity ϵ of the wall (related to the light speed in the wall) is used to perform matched filtering and compute wavenumbers κ_x and κ_z . However, we cannot count on a SiWa user to move his/her hand at a pre-determined speed. In addition, the diversity of wall materials can lead to very different permittivities. To make things worse, environment conditions such as humidity may also alter the properties of a wall drastically (e.g., the permittivity may vary from 5 to 14 at different moisture levels [56]). If these parameters are not specified in advance, existing algorithms cannot perform imaging successfully.

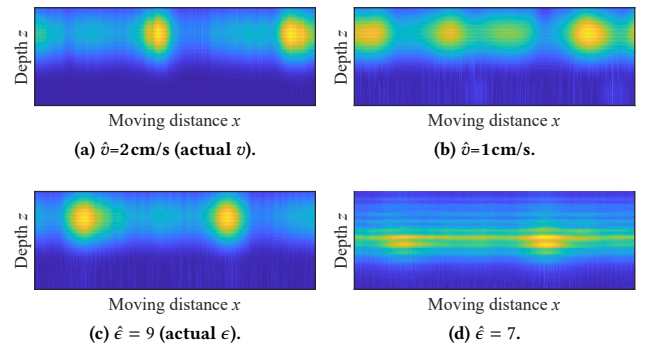


Figure 8: RMA imaging results affected by different parameter settings. (a) and (b) show the effect of the (radar) probe swiping speed v . (c) and (d) show the effect of relative permittivity ϵ of the wall.

To illustrate this difficulty, we use RMA for imaging two rebars in a concrete wall, whose actual permittivity ϵ is 9, and the speed of swiping the probe is about 2 cm/s. The results shown in Figure 8 clearly demonstrate that, if correct parameters are provided, the energy distribution is concentrated and the rebar boundaries are sharp. However, if the estimated velocities \hat{v} and permittivity $\hat{\epsilon}$ deviate from the actual values, the imaging results become blurred and may even deliver the wrong location information. As these weaknesses are inherent to existing model-based methods, it is imperative to equip SiWa with a calibration-free imaging module.

3.2.2 Material Identification. Suppose a map of the in-wall structure has been obtained by imaging, $\mathbf{r}(x, t)$ may further enable material identification and hence structural diagnosis. Two conventional methods for material identification are spectroscopy and ellipsometry, which respectively exploit dispersion and polarization.

Dispersion is the phenomenon of frequency differentiation of propagation speed. Because different materials exhibit distinct dispersion properties, this phenomenon can be leveraged to characterize respective materials via *spectroscopy*. The effect of dispersion manifests both in the time domain (broadening or narrowing of pulses) and frequency domain (skewed spectra). As a result, time-domain [18] and frequency-domain spectroscopy [32] have been developed to identify materials. In our case, the wall materials (e.g., concrete), subsurface structure (e.g., rebar), and seeping substances (e.g., water) together yield rich dispersion information. As an illustration, we probe a rebar and PVC pipe buried under concrete; the results in Figure 9 clearly showcase the distinct dispersion characteristics of the two materials.

We further study *polarimetry* (a.k.a. *ellipsometry*) founded on polarization [4, 46, 71, 85]. Essentially, when electromagnetic waves are reflected from the boundary of a target (of a certain material), their polarization is altered by the target's properties (e.g., refractive index, thickness, and dielectric function tensor) [4]. Therefore, if we deploy two antennas with linear orthogonal polarization, the readings from them should indicate the polarization rotation of the reflected waves, which in turn indicates the reflecting material. Normally, *reflection coefficients* γ_p and γ_s in the co-polarization and cross-polarization directions are adopted by polarimetry as features, they are represented as [48]:

$$\gamma_p = \frac{N_s \cos(\phi_0) - N_0 \cos(\phi_s)}{N_s \cos(\phi_0) + N_0 \cos(\phi_s)}, \gamma_s = \frac{N_0 \cos(\phi_0) - N_s \cos(\phi_s)}{N_0 \cos(\phi_0) + N_s \cos(\phi_s)}, \quad (5)$$

where N_0 and N_s are the refractive indices of the wall and the target material, and ϕ_0 and ϕ_s are the complex incidence and refraction angles determined by $\frac{\sin \phi_0}{\sin \phi_s} = \frac{N_s}{N_0}$ [35]. By measuring the reflection




Figure 9: Dispersions caused by different materials.

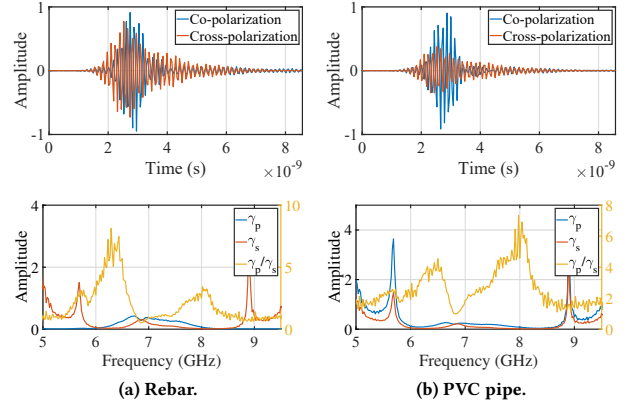


Figure 10: Reflected signals and reflection coefficients from orthogonally polarized antennas for different materials.

coefficients γ_p and γ_s , the refractive indices N_s of the target material can be inferred and in turn used to characterize the material.

We briefly validate the sensitivity of reflection polarization to materials. We record reflected signals from two orthogonally polarized antennas (see Figure 15) for rebar and PVC pipe buried in a concrete wall, and we also compute the reflection coefficients at all frequencies by dividing the spectrum of the received signal by that of the transmitted Gaussian pulse. As shown in Figure 10, distinct materials indeed cause different patterns in the waveforms. In accordance with Eq. (5), we also observe different ratios of γ_p to γ_s , evidently confirming the distinct rotations of reflected waves by two materials. We refer to [69] for a more detailed benchmark study on RF polarimetry for material identification.

Unfortunately, these promising techniques face two major issues for our application. On one hand, the signal difference may be conspicuous for two distinct materials, but what features to be used for differentiating a variety of materials is a tricky question. On the other hand, the above testing subjects are put at the same depth, but such ideal conditions do not exist in practice, hence environment influence such as wall materials and depth may greatly affect the reliability of these model-based techniques. Therefore, one has to identify novel methods that single out the relevant features automatically while masking the environment influence.

3.3 SiWa: A Deep Learning Approach

According to Section 3.2, model-based methods do not offer usable adaptability for our applications, although the signal matrix $\mathbf{r}(x, t)$ does contain sufficient information. To this end, we propose to apply a deep learning framework to integrate structural imaging and material identification functions while combating the challenges faced by model-based approaches. We refine the SiWa workflow in Figure 11, in which the hardware front-end remains (as in Figure 6), but $\mathbf{r}(x, t)$ is fed into the Imaging Network (I-Net) for producing structural imaging, and the output of I-Net guides the further processing of 1-D signal samples via Material identification Network (M-Net) for material identification and diagnosis.

3.3.1 SiWa Imaging Network (I-Net). Let us first recapitulate the limitations of SAR imaging algorithms: the speed of hand movement




Figure 11: A deep learning view of SiWa's design.

and the permittivity of a wall have to be specified beforehand; otherwise, the imaging results can be blurred and excessive noise can be caused by the wall, as shown in Figure 8. Fortunately, it is well-known that a *deep neural network* (DNN) can approximate arbitrary functions [13] and it can be trained to be robust against parameter changes [57]. Therefore, it makes sense for us to replace SAR imaging with a DNN that approximates the SAR capability while being adaptive to for example, wall properties. In fact, existing studies have already confirmed that commonly used filters can be approximated by DNNs [14, 22]. Moreover, recent advances in optics [37, 80] also support our idea that deep learning can be used for imaging purposes free of parameter tuning.

Specifically, we aim to design an encoder-decoder model [67] for “focusing” the RF signal matrix $\mathbf{r}(x, t)$ (in moving-range domain) into an image of the subsurface structure $\tilde{\mathbf{r}}(x, z)$ (in moving-depth domain). The encoder-decoder module consists of two parts: encoder ϕ and decoder ψ , with respective parameters θ_ϕ and θ_ψ . The former maps the input space \mathcal{R} to the (hidden) feature representation space \mathcal{H} , while the latter leverages \mathcal{H} to generate the output space $\tilde{\mathcal{R}}$. The training of I-Net is implemented by minimizing a loss function as follows:

$$\hat{\theta}_\phi, \hat{\theta}_\psi = \arg \min_{\theta_\phi, \theta_\psi} \|\hat{\mathbf{r}} - \psi(\phi(\mathbf{r}))\|_1; \quad (6)$$

the loss function measures the difference between a ground truth image (produced by RMA with correct parameters) $\hat{\mathbf{r}} \in \mathbb{R}^{N_x \times N_z}$ and an output image $\tilde{\mathbf{r}} = \psi(\phi(\mathbf{r}))$, where $\mathbf{r} \in \mathbb{R}^{N_x \times N_t}$, and N_x , N_z , and N_t are the number of samples along the moving, depth, and range axes, respectively.

The architecture of I-Net is shown in Figure 12. First, $\mathbf{r}(x, t)$ is fed into the multi-level encoder ϕ , a contraction network adopting regular Convolutional Neural Network (CNN) [43]. The feature representations distilled at every level of ϕ are subsequently fed to the peer level in the decoder ψ , which reverses the contraction and translates the feature representations to the imaging results. The decoder adopts more powerful ResNet blocks [29]. Replacing

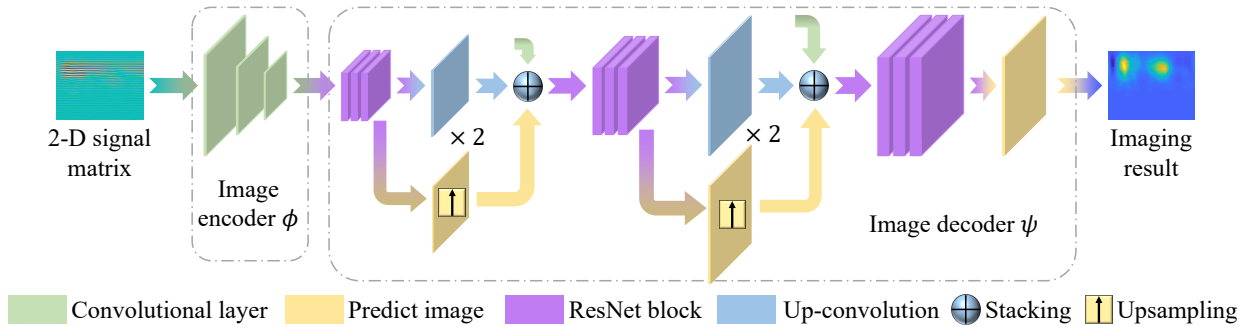


Figure 12: The encoder-decoder architecture of I-Net; the details are shown only for the decoder part.

pooling operators with upsampling operators, the resolutions of intermediate images are gradually increased. At each level of ψ , the features from i) skip connection with the peer level in ϕ (green), ii) up-convolutional layer (blue), and iii) upsampling layer (yellow) are stacked together and forwarded to the next level for further processing. Multiple levels of these functional blocks are connected in series to achieve a deeper network for better performance. As the last stage, a ResNet block followed by a convolutional layer gathers the intermediate image and generates the final result $\tilde{\mathbf{r}}(x, z)$. The sufficient capacity provided by the encoder-decoder model enables I-Net to discover environment parameters implicitly via ϕ , rendering structural imaging robust to changes of multiple environment conditions. With the structural imaging results provided at the output of I-Net, SiWa can now locate subsurface structures and further conduct material identification and diagnosis.

3.3.2 SiWa Material Identification Network (M-Net). In Section 3.2, we have explored the possibilities of employing the dispersion and polarization of the reflected signals for material identification. We also point out that these model-based methods face major difficulties in practice given the environment influence. To tackle this challenge, we aim to design M-Net as an effective classifier robust to environment influence. SiWa obtains the position of subsurface structures using the Constant False Alarm Rate (CFAR) algorithm [44] from the images produced by I-Net. It then extracts the 1-D signal sequence $r(t) \in \mathbb{R}^{N_t \times 1}$ from the 2-D radar signal matrix $\mathbf{r}(x, t) \in \mathbb{R}^{N_x \times N_t}$. Ideally, SiWa could perform beamforming to synthesize multiple $r(t)$ into one with stronger SNR, but that would require estimating the radar swiping speed v . Therefore, our current prototype takes the highest SNR $r(t)$ among those close to the known location.

As a classifier, M-Net G maps the input 1-D signal space \mathcal{R}^1 to the label space \mathcal{Y} that is a finite set indicating materials of different conditions (e.g., leaked PVC pipe, non-corroded rebar, etc.). Ideally, the parameters θ_G of M-Net can be trained by minimizing the cross-entropy loss $L(y, G(r))$ between the estimated and actual label:

$$\hat{\theta}_G = \arg \min_{\theta_G} L(y, G(r)). \quad (7)$$

However, the above formulation neglects the environment influence: it works only when the training and testing environments remain the same, because the training may overfit the features offered by the environment (e.g., wall thickness). To remove the environment influence on the classification results, we leverage

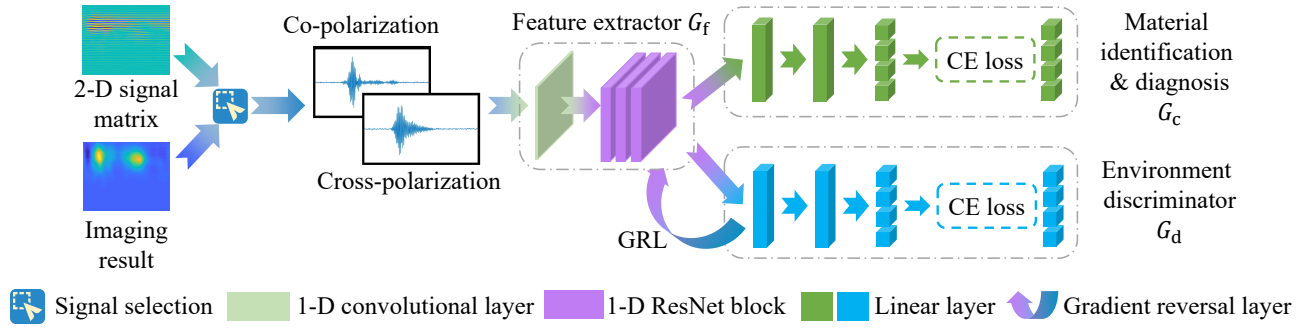


Figure 13: The architecture of M-Net along with its adversarial learning strategy.

the power of adversarial learning [21], and instead train M-Net as follows:

$$(\hat{\theta}_f, \hat{\theta}_c) = \arg \min_{\theta_f, \theta_c} L(y, e, r), \quad \hat{\theta}_d = \arg \max_{\theta_d} L(y, e, r), \quad (8)$$

where $L(y, e, r) = L_c(y, G_c(G_f(r))) - \lambda L_d(e, G_d(G_f(r)))$, G_f is the feature extraction network, G_c is the material classifier, and G_d is the environment discriminator; their corresponding parameters are θ_f , θ_c , and θ_d . L_c and L_d are cross-entropy losses for the classifier and discriminator, with λ controlling the trade-off between them.

The rationale behind this formulation is that, as M-Net is trained to predict material class labels using samples collected from different environments, the features used for classification should come only from materials but not from environments. To this end, an environment discriminator G_d is introduced to predict environment class labels e . While the training process aims to improve the environment classification accuracy by optimizing $\hat{\theta}_d$, an adversarial learning method is also taken to “cheat” G_d by negating the loss function via $-\lambda$. This procedure suppresses material-unrelated features and thus prevents M-Net from overfitting to specific environments. The network and its training strategy are illustrated in Figure 13. The feature extractor G_f takes in two data sequences with different polarization directions, and the extracted features are fed into the material classifier G_c and environment discriminator G_d . The loss negation via $-\lambda$ is achieved by reversing the gradient during the backpropagation process with the Gradient Reversal Layer (GRL), which only works for training purposes. In order to prevent conflicting gradients [83], we augment the backpropagation by projecting the gradients of G_f and G_c onto each other.

4 IMPLEMENTATIONS

Hardware Implementations. SiWa performs imaging and material identification using UWB signals. The core component is a compact and low-cost Novelda X4M05 [55] IR-UWB transceiver. It transmits a baseband signal with a bandwidth of 1.5GHz, and the signal is further modulated onto a 7.29GHz carrier. A Raspberry Pi single-board computer [63] is used to control the transceiver and interface with a laptop computer. The antennas are critical hardware components and they should have a good coupling ability with walls so as to deliver the maximum amount of power P_{in} into walls and minimize P_{ref} reflected by the air-wall boundary. Given this requirement, we define the SNR of an antenna as $\frac{P_{in}}{P_{ref}}$ to measure the performance of



Figure 14: Four antennas from left to right: log-periodic, Vivaldi, patch, and planar dipole antenna.

the antenna. We choose four common antenna types: log-periodic, Vivaldi, patch, and planar dipole antennas, as shown in Figure 14, and their respective SNRs are -4.17 dB, -17.23 dB, -22.21 dB, and -7.09 dB. Therefore, we choose log-periodic antennas for SiWa as they deliver the highest SNR. To utilize the information carried by reflected signals of different polarization states, as described in Section 3.2.2, we design a probe with 1 tx antenna, and 2 perpendicularly placed rx antennas (connected via a 2-port switch). The antennas are fixed into a foam block and spaced 4cm (about 1 wavelength) apart to avoid mutual coupling. The whole hardware prototype is shown in Figure 15.

Software Implementations. SiWa’s imaging and material identification modules are implemented using Python 3.7 and PyTorch 1.7.1, and trained using GeForce RTX 2080 Ti graphics card. CPUs from PC workstations and single-board computers are used during the inference stage of structural imaging and material identification. The software module for controlling the X4 transceiver and antennas is implemented in C++. The frame rate of the X4 transceiver is set to 40 frames per second. To get the most out of the X4 chip,




Figure 15: Hardware prototype of SiWa.




Figure 16: We perform extensive field tests in a precast concrete plant. Concrete blocks with different shapes, mix ratios, water contents, and aging degrees are investigated under our tests.

we modify its driver program to directly obtain the raw RF signal, instead of the downsampled baseband signal accessible by default.

The parameters of I-Net are set as follows: for the 2-D convolutional layer, we set the kernel size to 3, stride to 1, padding size to 1, and the number of output channels is 2 times the number of input channels; for the 2-D up-convolutional layer, we set the kernel size to 3, stride to 1, scale factor to 2, and the number of input channels is 2 times of the number of output channels. All weights are initialized by the Xavier uniform initializer [41]. We divide the collected dataset into training and test sets, based on the principle that they come from different environments. The training set contains 8,000 pairs of 2-D signal matrix and imaging ground truth, and the test set contains 6,000 pairs. The size of a training image is 200×200 (though I-Net can handle variable-length inputs during inference), and the ground truth images are obtained by RMA. As a preprocessing, the 2-D radar signal matrices are normalized to have zero mean and a standard deviation of 0.017, the ground truth images are normalized to a mean of 0.092 and standard deviation of 0.2423. For the training process, we set the batch size to 64, adopt the L^1 loss, and use the Stochastic Gradient Descent optimizer [6], whose learning rate and momentum are set to 0.001 and 0.9, respectively.

The parameters of M-Net are set as follows: the kernel size of the convolutional layer is 3, and the stride is 2, λ that controls the trade-off between G_c and G_d is 1. G_d consists of 3 consecutive fully-connected layers, transforming the 15,360 features output by G_f to the final layer representing different environments. We divide the collected dataset to a training set containing 8,000 1-D signals collected directly and a test set contains 6,000 1-D signals extracted from 2-D signal matrices, and the training and test sets are so partitioned that they come from different environments. The length of each signal sequence is 1120. In the preprocessing stage, the two-channel signals are normalized to a mean of (0, 0) and a standard deviation of (0.0052, 0.0021). During the training process, we set the batch size to 128, adopt cross-entropy loss for both the material classifier and environment discriminator, and use the Adam optimizer [40] with a learning rate of 0.001.

5 EVALUATION

In this section, we evaluate the performance of SiWa. We embed different objects (corroded rebar, non-corroded rebar, PVC pipe, and leaked PVC pipe) in various walls. The rebar is HRB400-grade and has a nominal diameter of 10mm, and the PVC pipe has a diameter of 20mm and a wall thickness of 4mm. Aside from common drywall (made of gypsum), brick, and wooden walls that can be readily handled in our lab, we pay special attention to concrete walls since

they are the most important building material, yet they cause the highest loss to RF signals. We have conducted extensive field tests on concrete blocks with different shapes, mix ratios, water contents, aging degrees, and embedded materials in a precast concrete plant, as shown in Figure 16; the ground truth positions, depths, materials of the embedded objects are all clearly labeled.

In the experiment, the permittivity of a wall is measured using a Keysight N5234B PNA-L vector network analyzer [39]. We maintain a relatively stable hand moving speed by following the visual marks on the wall and the auditory beats generated by a metronome. The hand moving speed is then derived from distance changes per unit time measured by a Bosch GLM 50C laser measure [65]. The outcome is that every collected signal matrix is induced by a different nominal speed along with minor (unknown) speed variations. We shall refrain from comparing I-Net with through-wall [3, 38] and in-air [3, 24, 38, 90] RF imaging techniques, because i) through-wall Wi-Fi sensing fails to support in-wall imaging, as proven by Figure 4a, and ii) in-air imaging faces challenges such as excessive air-wall boundary loss and varying environment parameters, as demonstrated in Figures 5 and 8. Proposals for material identification (e.g., [23]) are also excluded from comparing with M-Net, as they often classify materials based on photos and are thus totally irrelevant to RF-sensing. In the following, we specify comparison baselines for structural imaging and material identification and report their evaluation results respectively.

5.1 Structural Imaging

We start with evaluating the performance of the I-Net proposed in Section 3.3.1; a qualitative result is provided first for a concrete block with rebars embedded, where we mark the ground truth labels in Figure 17a and report the imaging results obtained by swiping the probe for a 40cm distance along the concrete surface in Figure 17b. We can see that SiWa accurately produces the structural image within the concrete. Note that what we see are the cross-sections of rebars; multiple parallel scans are needed to obtain a 3-D tomography map that extends along their axial direction. We hereafter report more on I-Net's quantitative evaluations. Moreover, we compare I-Net with RMA [7] instead of other methods summarized in [90], as it is the most practiced in-air SAR imaging algorithm capable of obtaining high resolution if all parameters are properly specified. We also refrain from comparing with Walabot [78], because it is a closed-source system that cannot be customized.

5.1.1 Horizontal Distance Accuracy. Horizontal distance accuracy along the radar scanning direction x is first evaluated. The error

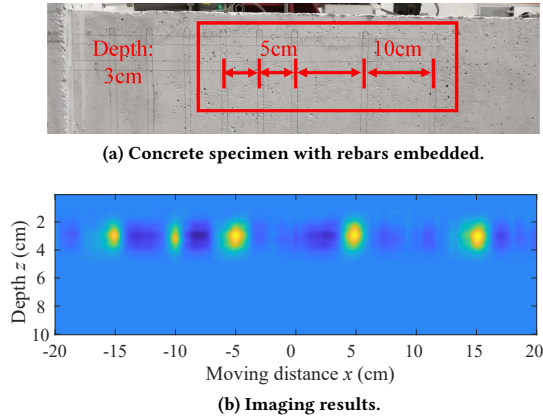


Figure 17: A qualitative test on a concrete block with five rebars embedded.

is defined as the absolute difference between the estimated horizontal distance x_e and the actual horizontal distance x_a , namely, $|x_e - x_a|$. For RMA, we consider two cases: i) the parameters (i.e., the permittivity of a wall and the speed of hand movement) of the algorithms are accurate, and ii) the parameters have errors within $\pm 20\%$, representing common parameter estimation situations in real-life usage. Figure 18 plots the Cumulative Distribution Functions (CDFs) of horizontal distance error given different walls. For I-Net, we observe from Figure 18a that the median errors are 0.2cm, 0.4cm, 0.7cm, and 0.7cm for drywall, wooden wall, brick wall, and concrete wall, respectively. These errors are in line with common sense that concrete and brick walls cause higher attenuation to RF signals, hence larger localization errors. As a comparison, RMA, whose CDFs shown in Figure 18b, has slightly better performance: median errors 0.2 cm, 0.2 cm, 0.3 cm, and 0.4 cm given the same four cases, if the parameters are set correctly. However, if the estimated parameters deviate from the actual values, the median errors increase to 0.3cm, 0.4cm, 0.5cm, and 3.6cm (legend bearing asterisk in the plot), and the curves (dashed lines) have very long tails, indicating a serious performance degradation.




Figure 18: CDFs of horizontal distance errors.

5.1.2 Depth Accuracy. Next, we evaluate the performance of SiWa with respect to the depth of subsurface structures. The error is defined as the absolute difference between the estimated depth z_e and the actual depth z_a , namely, $|z_e - z_a|$. We again compare I-Net with RMA, based on settings similar to Section 5.1.1. Figure 19a shows I-Net's median errors of 0cm, 0.1cm, 0.2cm, and 0.4cm for wooden wall, drywall, brick wall, and concrete wall, respectively; the same intuition on horizontal distance applies here too. Note that




Figure 19: CDFs of depth errors.

the non-smoothness of the CDF curve is caused by the limited range resolution of a microwave radar, yet mmWave radars delivering a higher range resolution fail to work (see Section 2.2.3). RMA (see Figure 19b) again leads to negligible errors given correct parameters, but non-accurate parameters cause it to deliver totally wrong results with median error up to 3cm, definitely not suitable for practical adoption. The results in Sections 5.1.1 and 5.1.2 clearly indicate that I-Net is robust to multi-parameter changes.

5.1.3 Imaging Quality. Since the imaging results will be used to assist inspectors to locate subsurface structures, it should be clear and visually appealing. Figure 20 illustrates the imaging quality difference between I-Net and RMA with correct parameters. The I-Net imaging result is smoother and mostly free of noise, while the result of RMA is noisy due to the existence of high-frequency components. To measure the imaging quality quantitatively, we cannot use conventional SNR, since we do not have a reference image. Therefore, we adopt the Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE) [53] score. The BRISQUE method compares an imaging result to a default model computed from images of natural scenes with similar distortions; it yields a unitless score between 0 and 100, with a smaller score value indicating a better perceptual quality. The BRISQUE scores of I-Net and RMA are shown in Table 1, confirming a slight advantage of I-Net over RMA in terms of overall image quality given all wall materials.




Figure 20: Imaging results of two algorithms. The result of I-Net is smoother than that of RMA.

Table 1: BRISQUE scores of I-Net and RMA.

	Concrete	Brick	Wooden	Drywall
I-Net	40.24	45.28	34.80	38.57
RMA	45.44	49.34	40.15	40.84

5.1.4 Impact of Water Content. We then evaluate the effect of water content on the imaging result. In many countries with rainy and humid climates, the water content in the wall can vary greatly and affect the performance of SiWa. This effect on imaging is twofold: i) water absorbs electromagnetic radiation in the RF spectrum ranging

from 1 GHz to 1000 GHz [28, 79], and ii) water causes a drastic permittivity change of a wall [56], both leading to the failure of RMA. The evaluation results in Figure 21a show that as the water content in wall increases from 0 g/m^3 to 100 g/m^3 , the median horizontal distance error goes up for 0.3cm on average, while the median depth error increases by around 0.2cm. The gradual increase of error indicates that, as the water content increases, more RF signals are absorbed/blocked. It is worth noting that the increase in water content does not lead to a steep rise in error, confirming the robustness of I-Net in adapting to permittivity variations and validating our claim that I-Net is calibration-free.

5.1.5 Impact of Depth. RF signals are greatly attenuated as they travel deep into walls, hence the depth of an in-wall structure is a key limiting factor of SiWa's imaging and localization capabilities. We now test whether I-Net can discover subsurface structures deep inside the wall. We put different materials at depths of 3 cm, 5 cm, and 10 cm, and evaluate the accuracies in both horizontal distance and depth, with the results shown in Figure 21b. Take concrete wall as an example, the median horizontal distance errors are 0.4 cm, 0.7 cm, and 1.5 cm while the depth errors are 0.3 cm, 0.5 cm, and 1.2 cm, respectively. Other walls follow the same trend but appear less sensitive to depth. Apparently, SiWa is still able to deliver a decent accuracy as the depth increases.

5.1.6 Impact of Hand Movement Speed. As we discussed in Section 3.2, traditional RMA algorithm is sensitive to the movement speed of the radar probe. In this section, we validate that the I-Net used in SiWa is not sensitive to this speed. We let a SiWa user move the probe on a wall with constant speeds of 1cm/s, 2cm/s, and 4cm/s. The evaluation results are shown in Figure 21c. It can be observed that the median errors in horizontal distance and depth remain almost the same. This confirms that I-Net is robust to different hand movement speeds. We do observe that the error variances slightly grow as the speed increases, which may partially be attributed to the fact that the SiWa user cannot swipe the probe steadily on the wall at a higher speed. As we explained earlier, a distinct speed variation pattern happens to every trial regardless of how hard a user tries to maintain a stable speed. Fortunately, I-Net is shown to be robust to such minor variations.

5.1.7 Impact of Material. The material of subsurface structures also plays an important role in the imaging and localization of I-Net. To investigate the effects of different materials, we choose three common ones: rebar, water-filled PVC pipe, and empty PVC pipe. In Figure 21d, we observe that, for rebar and water-filled PVC pipes, more than 75 % of the horizontal distance and depth errors are smaller than 1 cm; although empty PVC pipes incur slightly higher errors in percentile, all median errors in horizontal distance and depth are below 1 cm. We believe that the difference in reflectivity (metal and water are strong reflectors, but PVC plastics are weaker) may have caused the discrepancy in errors.

5.1.8 I-Net Efficiency. As the last step of evaluating I-Net, we investigate its runtime latency. In order to make our results compatible with most computers, we consider CPU-only inference, and we compare our algorithm with RMA on different devices. The PC workstation we use is equipped with Intel Xeon W-2133 CPU and 16 GB RAM, the laptop is equipped with AMD Ryzen 5 3500U CPU and 8 GB RAM, Raspberry Pi 3 is equipped with Arm Cortex-A53 CPU and 1 GB RAM, and Raspberry Pi 4 is equipped with ARM Cortex-A72 CPU and 2 GB RAM. Assuming an input 2-D signal matrix of size 200×200 , Table 2 reports the inference times under the above four devices. It appears that the runtime latency of I-Net is slightly higher than that of RMA, but the latency of I-Net on PC workstation and laptop can satisfy the requirement of real-time processing. We plan to further optimize the structure of I-Net, so as to reduce its complexity for embedded platforms (e.g., Raspberry Pi) while maintaining the imaging performance. This should enable a standalone SiWa prototype with real-time capability.

Table 2: Runtime latency of RMA and I-Net.

	PC Workstation	Laptop	RPi 3	RPi 4
RMA	0.05 s	0.12 s	2.97 s	1.22 s
I-Net	0.08 s	0.15 s	3.09 s	1.43 s

5.2 Material Identification

In this section, we first discuss the overall material identification performance of SiWa's M-Net. Then we investigate M-Net's robustness to depth and water content changes. We also discuss the

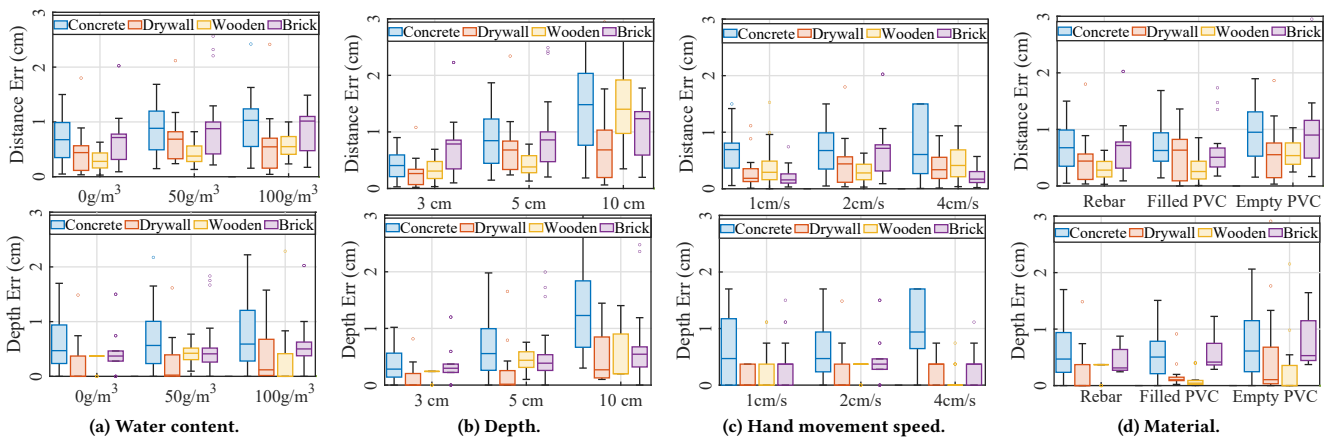


Figure 21: Relationship between localization errors and various environment and operating variables.

necessity of dual polarization for accurate material identification. We further compare M-Net to 2-D ResNet [29] and Multi-Layer Perceptron (MLP) [59]. ResNet is a popular deep learning network for classification tasks, and MLP is a common feed-forward network used in a recent proposal [85] for material identification. At last, we verify the computation efficiency of M-Net.

5.2.1 Overall Performance. We prepare four different in-wall structures: non-corroded rebar, corroded rebar, non-leaked PVC pipe, and leaked PVC pipe. We first use confusion matrices to characterize the identification performance under different walls in Figure 22, where we observe that concrete and brick walls have more adverse effects on material identification than wooden wall and drywall. It is also interesting to note the correlated adverse effect between concrete and rebar, as well as that between brick and PVC. By summarizing these data, we could also derive that the overall accuracy for correctly identifying material (i.e., steel or PVC) is 95.2%, while that for further correctly diagnosing their health status is slightly lower at 91.8%.

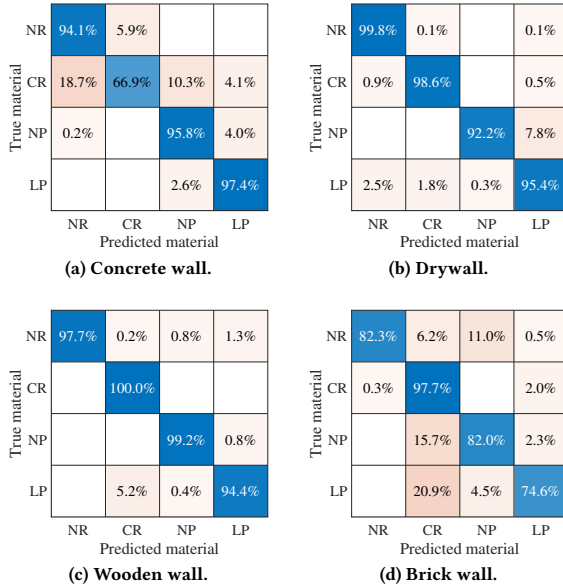


Figure 22: Confusion matrices of M-Net under different walls, where non-corroded rebar, corroded rebar, non-leaked PVC pipe, and leaked PVC pipe are abbreviated as NR, CR, NP, LP, respectively.

5.2.2 Depth. As discussed in Section 5.1.5, attenuation of RF signals as they travel into walls makes it harder to detect in-wall structures. Therefore, we hereby study the effects of depth on material identification in this section, resulting in Figure 23 to depict that the accuracy slightly decreases as the depth increases. Fortunately, while the identification accuracies are all over 90%, the diagnosis accuracies remain above 80%. The worst performance caused by brick walls may result from their impurity (e.g., cavities or stones inside), but these are also realities that have to be faced by SiWa. Overall, these results validate that our adversarial training strategy successfully removes the influence caused by varying wall depths and makes M-Net virtually immune to environment changes.




Figure 23: Impact of depth on material identification.

5.2.3 Water Content. We then test M-Net's robustness to various water contents in the wall. We gradually spray water onto the wall, from 0g/m^3 to 100g/m^3 , and employ M-Net to identify and diagnose in-wall materials. According to the results shown in Figure 24, we observe that, though performance degrades due to water blocking and absorbing RF signals, the overall identification accuracies and diagnosis accuracies remain over 90% and 80%, respectively. These results again confirm the robustness of M-Net.




Figure 24: Impact of water on material identification.

5.2.4 Polarization Diversity. In Section 3.2.2, we have discussed polarimetry exploiting the diversity of polarizations to identify materials. To prove the necessity of using two polarizations, we perform an ablation study by removing the cross-polarization. The overall identification and diagnosis accuracies are 95.2% and 91.8% for dual polarization, but 86.1% and 70.5% for single polarization. These large gaps validate the need for dual polarization.

5.2.5 Other Neural Networks. Table 3 summarizes the material identification performance with different deep neural networks, where M-Net* is a variant of M-Net without environment discriminator. To adapt 2-D ResNet for processing radar signal, we first transform the 1-D signal sequence to 2-D spectrogram containing both temporal and frequency information, then we employ the network to classify the spectrogram. Clearly, M-Net achieves the highest accuracy while maintaining a reasonable processing latency. We also observe that the runtime latency of M-Net on single-board Raspberry Pi can potentially satisfy real-time requirements.

Table 3: Comparing M-Net with other networks.

	M-Net	M-Net*	2-D ResNet	MLP [85]
Identification acc (%)	95.2	92.9	91.4	89.8
Diagnosis acc (%)	91.8	85.1	86.4	78.9
Latency on PC (s)	0.032	0.031	0.057	0.013
Latency on RPi 4 (s)	0.96	0.96	1.63	0.51

5.3 Case Studies

In this section, we extend the aforementioned evaluations to more complex scenarios, including cracks in wall and multi-material wall. Due to the lack of sufficient specimens, we do not re-train SiWa. Instead, we examine if the structural imaging ability of I-Net trained on previous samples can be generalized to these complex scenarios.

5.3.1 Cracks in Wall. Wall crack detection [20] is a critical goal of non-destructive inspection, and it is hence an objective of SiWa too. To verify SiWa's ability in crack detection, we prepare a cracked concrete block, as shown in Figure 25a, and we apply another cement layer on top of the cracked surface. This allows us to emulate in-wall cracks while still having full knowledge of the ground truth. Then we scan this augmented sample to test the imaging performance of SiWa in visualizing these cracks.




Figure 25: Three scanning strips (with ground truth cracks marked) and their corresponding imaging results.

We scan three horizontal strips (marked along with underlying cracks in Figure 25a) on the newly applied cement layer; the resulting SiWa imaging outcomes are depicted in Figure 25b, with each subfigure showing a 2-D cross-section image along the width x (moving dimension) and depth z . In these images, one may evidently discern major cracks and dents (marked red in Figure 25a) as bright spots. One may also observe that small cracks only generate bright patches and hence may not be precisely identified. This phenomenon can be attributed to the relatively weak reflection and the lack of crack samples in the training dataset. Nonetheless, the current results clearly demonstrate that SiWa is sufficiently competent to identify major issues even without any training.

5.3.2 Multi-material Wall. By far we have considered only homogeneous walls made of a single material, but multi-layer walls made of different materials can be adopted for interior finish; for example, wooden tiles are often used to cover concrete walls/floors for decoration purpose. To prove that SiWa's in-wall detection ability can be generalized to multi-material walls, we cover a reinforced concrete block with 1.5 cm thick wooden tile, as shown in Figure 26a. We can see that the two rebars are approximately 4.5 cm deep and 6 cm apart in the moving distance x direction.

We then conduct SiWa scan on the wooden surface, and the SiWa imaging result is shown in Figure 26b. One may readily observe the two rebars imaged with high localization accuracy in both directions of horizontal distance x and depth z . One may also notice that the wood-concrete interface introduces a certain level of noise to the imaging result and causes the imaging quality to degrade slightly. This can be explained by the fact that the signals are partly




Figure 26: Multi-material wall and imaging result.

reflected back at the material discontinuity, hence negatively impacting SNR. Nevertheless, the current result is sufficiently good to confirm that SiWa supports imaging into multi-material walls, and we believe better imaging quality can be achieved if a sufficient amount of multi-material wall samples become available.

6 LIMITATIONS AND FUTURE WORK

As SiWa is a proof-of-concept prototype using IR-UWB radar and novel deep learning modules for seeing into walls, it may still leave a number of technical limitations for future exploration. First, we only report 2-D imaging in Section 5, as achieving tomography-like 3-D imaging is just a straightforward extension by multiple parallel scans (which is a common practice). Therefore, we consider implementing an antenna array for more effective and efficient 3-D imaging as a future direction [10]. Second, as a follow-up of the previous point, we also need to consider using deep generative models [89] to improve the robustness of the imaging ability (especially for 3-D imaging). Third, we have only tested four different objects in the material identification experiments, but we believe that M-Net can identify other materials as well. We are discussing with local industries to put SiWa into use, so that we may gather more training and performance indicating data. Last but not least, as the latency of I-Net and M-Net on single-board computers are only quasi-real-time, we are planning to use techniques such as model compression [5, 49, 82] to speed them up.

7 CONCLUSION

Seeing into walls without undermining structural integrity is a challenging yet important problem pertinent to building health diagnostics. To this end, we have implemented SiWa, a portable and low-cost in-wall scanning prototype based on an IR-UWB radar, aiming to synthesize in-wall images and identify the material status. Employing carefully designed deep learning modules, SiWa performs structural imaging and material identification robustly in a self-adaptive manner. With extensive experiments on different walls and materials, we have demonstrated the promising performance of SiWa in accurate imaging and effective identification robust to various environments. We are on the way to identifying practical deployment scenarios for performing real-life tests on SiWa, so as to push it to the application market.

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers for their valuable and constructive comments. We further thank Energy Research Institute @ NTU (ERI@N) and Interdisciplinary Graduate Programme (IGP) of NTU for supporting the PhD scholarship of Tianyue Zheng.

REFERENCES

- [1] ACS-Solutions GmbH. 2021. A1220 MONOLITH 3D. <https://acs-international.com/product/a1220-monolith-classic/>. Accessed: 2021-02-22.
- [2] Fadel Adib and Dina Katabi. 2013. See Through Walls with WiFi!. In *Proc. of ACM SIGCOMM*. 75–86.
- [3] Moeness G. Amin. 2017. *Through-the-Wall Radar Imaging*. CRC press.
- [4] Rasheed M.A. Azzam, Nicholas Mitchell Bashara, and Stanley S. Ballard. 1978. Ellipsometry and Polarized Light. *Physics Today* 31, 11 (1978), 72.
- [5] Sourav Bhattacharya and Nicholas D. Lane. 2016. Sparsification and Separation of Deep Learning Layers for Constrained Resource Inference on Wearables. In *Proc. of the 14th ACM SenSys*. 176–189.
- [6] Léon Bottou. 2012. Stochastic Gradient Descent Tricks. In *Neural networks: Tricks of the trade*. Springer, 421–436.
- [7] Ciro Cafforio, Claudio Prati, and Fabio Rocca. 1991. SAR Data Focusing using Seismic Migration Techniques. *IEEE Trans. Aerospace Electron. Systems* 27, 2 (1991), 194–207.
- [8] Yee Kit Chan and Voon Chet Koo. 2008. An Introduction to Synthetic Aperture Radar (SAR). *Progress In Electromagnetics Research* 2 (2008), 27–60.
- [9] Zhe Chen, Tianyue Zheng, Chao Cai, and Jun Luo. 2021. MoVi-Fi: Motion-robust Vital Signs Waveform Recovery via Deep Interpreted RF Sensing. In *Proc. of the 27th ACM MobiCom*. 1–14.
- [10] Zhe Chen, Tianyue Zheng, and Jun Luo. 2021. Octopus: A Practical and Versatile Wideband MIMO Sensing Platform. In *Proc. of the 27th ACM MobiCom*. 1–14.
- [11] Oliver Anthony Clark and Anthony Clark. 2003. *Seeing Beneath the Soil: Prospecting Methods in Archaeology*. Routledge.
- [12] Ned M. Cleland. 2008. Structural Integrity and Progressive Collapse in Large-Panel Precast Concrete Structural Systems. *PCI journal* 53, 4 (2008), 55.
- [13] Balázs Csanád Csáji et al. 2001. Approximation with Artificial Neural Networks. *Faculty of Sciences, Eötvös Loránd University, Hungary* 24, 48 (2001), 7.
- [14] Tri Dao, Albert Gu, Matthew Eichhorn, Atri Rudra, and Christopher Ré. 2019. Learning Fast Algorithms for Linear Transforms using Butterfly Factorizations. In *Proc. of the 11st ICLR*. 1517–1527.
- [15] Mita D. Desai and William Kenneth Jenkins. 1992. Convolution Backprojection Image Reconstruction for Spotlight Mode Synthetic Aperture Radar. *IEEE Transactions on Image Processing* 1, 4 (1992), 505–517.
- [16] Shuya Ding, Zhe Chen, Tianyue Zheng, and Jun Luo. 2020. RF-Net: A Unified Meta-Learning Framework for RF-Enabled One-Shot Human Activity Recognition. In *Proc. of the 18th ACM SenSys*. 517–530.
- [17] Massimo Donelli and Federico Viani. 2017. Remote Inspection of the Structural Integrity of Engineering Structures and Materials with Passive MST Probes. *IEEE Transactions on Geoscience and Remote Sensing* 55, 12 (2017), 6756–6766.
- [18] Timothy D. Dorney, Richard G. Baraniuk, and Daniel M. Mittleman. 2001. Material Parameter Estimation with Terahertz Time-Domain Spectroscopy. *JOSA A* 18, 7 (2001), 1562–1571.
- [19] Gustavo S. Duffó and Silvia B. Farina. 2009. Development of an Embeddable Sensor to Monitor the Corrosion Process of New and Existing Reinforced Concrete Structures. *Construction and Building Materials* 23, 8 (2009), 2746–2751.
- [20] Cao Vu Dung and Le Duc Anh. 2019. Autonomous Concrete Crack Detection using Deep Fully Convolutional Neural Network. *Automation in Construction* 99 (2019), 52–58.
- [21] Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised Domain Adaptation by Backpropagation. In *Proc. of the 32nd ICML*. 1180–1189.
- [22] Daniel George and Eliu Antonio Huerta. 2018. Deep Learning for Real-Time Gravitational Wave Detection and Parameter Estimation: Results with Advanced LIGO Data. *Physics Letters B* 778 (2018), 64–70.
- [23] Navid Ghassemi, Hadi Mahami, Mohammad Tayarani Darbandi, Afshin Shoebi, Sadiq Hussain, Farnad Nasirzadeh, Roohallah Alizadehsani, Darius Nahavandi, Abbas Khosravi, and Saeid Nahavandi. 2020. Material Recognition for Automated Progress Monitoring using Deep Learning Methods. *arXiv preprint arXiv:2006.16344* (2020).
- [24] Junfeng Guan, Sohrab Madani, Suraj Jog, Saurabh Gupta, and Haitham Hassanieh. 2020. Through Fog High-Resolution Imaging Using Millimeter Wave Radar. In *Proc. of the 33rd IEEE/CVF CVPR*. 11464–11473.
- [25] Udaya B. Halabe, Archana Vasudevan, Powsiri Klinkhachorn, and Hota V.S. GangaRao. 2007. Detection of Subsurface Defects in Fiber Reinforced Polymer Composite Bridge Decks using Digital Infrared Thermography. *Nondestructive Testing and Evaluation* 22, 2-3 (2007), 155–175.
- [26] J. Todd Halsey and Richard Miller. 1996. Destructive Testing of Two Forty-Year-Old Prestressed Concrete Bridge Beams. *PCI journal* 41, 5 (1996), 84–93.
- [27] Baoguo Han, Liqing Zhang, and Jinping Ou. 2017. *Smart and Multifunctional Concrete Toward Sustainable Infrastructures*. Springer.
- [28] J.B. Hasted and M.A. Shah. 1964. Microwave Absorption by Water in Building Materials. *British Journal of Applied Physics* 15, 7 (1964), 825.
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proc. of the 29th IEEE/CVF CVPR*. 770–778.
- [30] James Helal, Massoud Sofi, and Priyan Mendis. 2015. Non-Destructive Testing of Concrete: A Review of Methods. *Electronic Journal of Structural Engineering* 14, 1 (2015), 97–105.
- [31] Brian Hobbs and Mohamed Tchoketch Kebir. 2007. Non-Destructive Testing Techniques for the Forensic Engineering Investigation of Reinforced Concrete Buildings. *Forensic Science International* 167, 2-3 (2007), 167–172.
- [32] J. Michael Hollas. 2004. *Modern Spectroscopy*. John Wiley & Sons.
- [33] Graeme Hopkins and Christine Goodwin. 2011. *Living Architecture: Green Roofs and Walls*. Csiro Publishing.
- [34] Yoshinao Hoshi, Taisuke Koike, Hiroyuki Tokieda, Isao Shitanda, Masayuki Itagaki, and Yoshitaka Kato. 2019. Non-Contact Measurement to Detect Steel Rebar Corrosion in Reinforced Concrete by Electrochemical Impedance Spectroscopy. *Journal of The Electrochemical Society* 166, 11 (2019), 3316–3319.
- [35] Keigo Iizuka and Keigo Iizuka. 2008. *Engineering Optics*. Vol. 35. Springer.
- [36] Wenjun Jiang, Hongfei Xue, Chenglin Miao, Shiyang Wang, Sen Lin, Chong Tian, Srinivasan Murali, Haochen Hu, Zhi Sun, and Lu Su. 2020. Towards 3D Human Pose Construction Using WiFi. In *Proc. of the 26th ACM MobiCom*. 1–14.
- [37] Yuncheng Jin, Yiye Zhang, Lejia Hu, Haiyang Huang, Qiaoqi Xu, Xinpei Zhu, Limeng Huang, Yao Zheng, Hui-Liang Shen, Wei Gong, et al. 2018. Machine Learning Guided Rapid Focusing with Sensor-less Aberration Corrections. *Optics Express* 26, 23 (2018), 30162–30171.
- [38] Chitra R. Karanam and Yasamin Mostofi. 2017. 3D Through-Wall Imaging with Unmanned Aerial Vehicles Using WiFi. In *Proc. of the 16th ACM/IEEE IPSN*. 131–142.
- [39] Keysight Technologies. 2021. N5234B PNA-L Microwave Network Analyzer, 43.5 GHz. <https://www.keysight.com/en/pdx-2813612-pn-N5234B/pna-l-microwave-network-analyzer-435-ghz?cc=SG&lc=eng>. Accessed: 2021-03-02.
- [40] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [41] Siddharth Krishna Kumar. 2017. On Weight Initialization in Deep Neural Networks. *arXiv preprint arXiv:1704.08863* (2017).
- [42] Léonard Le Jeune, Sébastien Robert, Eduardo Lopez Villaverde, and Claire Prada. 2016. Plane Wave Imaging for Ultrasonic Non-Destructive Testing: Generalization to Multimodal Imaging. *Ultrasonics* 64 (2016), 128–138.
- [43] Yann LeCun, Patrick Haffner, Léon Bottou, and Yoshua Bengio. 1999. Object Recognition with Gradient-Based Learning. In *Shape, Contour and Grouping in Computer Vision*. Springer, 319–345.
- [44] Nadav Levanon. 1988. *Radar Principles*. New York (1988).
- [45] Kun Li, Bin Liang, Jing Yang, Jun Yang, and Jian-chun Cheng. 2017. Acoustic Broadband Metacouplers. *Applied Physics Letters* 110, 20 (2017), 1–5.
- [46] Tianxing Li, Derek Bai, Temiloluwa Prioleau, Nam Bui, Tam Vu, and Xia Zhou. 2020. Noninvasive Glucose Monitoring using Polarized Light. In *Proc. of the 18th ACM SenSys*. 544–557.
- [47] Tong Li, Nestore Galati, J. Gustavo Tumialan, and Antonio Nanni. 2005. Analysis of Unreinforced Masonry Concrete Walls Strengthened with Glass Fiber-Reinforced Polymer Bars. *ACI structural journal* 102, 4 (2005), 569.
- [48] John C. Lindon, George E. Tranter, and David Koppenaal. 2016. *Encyclopedia of Spectroscopy and Spectrometry*. Academic Press.
- [49] Sicong Liu, Yingyan Lin, Zimu Zhou, Kaiming Nan, Hui Liu, and Junzhao Du. 2018. On-Demand Deep Model Compression for Mobile Devices: A Usage-Driven Model Selection Framework. In *Proc. of the 16th ACM MobiSys*. 389–400.
- [50] Chris Xiaoxuan Lu, Stefano Rosa, Peijun Zhao, Bing Wang, Changhao Chen, John A Stankovic, Niki Trigoni, and Andrew Markham. 2020. See through Smoke: Robust Indoor Mapping with Low-Cost mmWave Radar. In *Proc. of the 18th ACM MobiSys*. 14–27.
- [51] Christiane Maierhofer, Hans-Wolf Reinhardt, and Gerd Dobmann. 2010. *Non-Destructive Evaluation of Reinforced Concrete Structures: Non-Destructive Testing Methods*. Elsevier.
- [52] Wenguang Mao, Mei Wang, and Lili Qiu. 2018. AIM: Acoustic Imaging on a Mobile. In *Proc. of the 16th ACM MobiSys*. 468–481.
- [53] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. 2012. No-Reference Image Quality Assessment in the Spatial Domain. *IEEE Transactions on Image Processing* 21, 12 (2012), 4695–4708.
- [54] David C. Munson, James D. O'Brien, and W. Kenneth Jenkins. 1983. A Tomographic Formulation of Spotlight-Mode Synthetic Aperture Radar. *Proc. IEEE* 71, 8 (1983), 917–925.
- [55] Novelda AS. 2017. The World Leader in Ultra Wideband (UWB) Sensing. <https://novelda.com/technology/>. Accessed: 2021-02-22.
- [56] Ade Ogunsoola, U. Reggiani, and Leonardo Sandrolini. 2005. Shielding Effectiveness of Concrete Buildings. In *IEEE 6th International Symposium on Electromagnetic Compatibility and Electromagnetic Ecology*, 2005. IEEE, 65–68.
- [57] Francesco Orabona. 2014. Simultaneous Model Selection and Optimization through Parameter-Free Stochastic Learning. *Proc. of the 27th NIPS* 2, January (2014), 1116–1124.
- [58] Caner Özdemir, Şevket Demirci, Enes Yiğit, and Betül Yılmaz. 2014. A Review on Migration Methods in B-Scan Ground Penetrating Radar Imaging. *Mathematical Problems in Engineering* 2014 (2014).
- [59] S.K. Pal and S. Mitra. 1992. Multilayer Perceptron, Fuzzy Sets, and Classification. *IEEE Transactions on Neural Networks* 3, 5 (1992), 683–697.
- [60] V.V. Panasyuk, V.I. Marukha, and V.P. Sylovanyuk. 2013. *Injection Technologies for the Repair of Damaged Concrete Structures*. Springer.

- [61] Howard W Pickering. 2003. Important Early Developments and Current Understanding of the IR Mechanism of Localized Corrosion. *Journal of the Electrochemical Society* 150, 5 (2003), 1–13.
- [62] Akarsh Prabhakara, Vaibhav Singh, Swarn Kumar, and Anthony Rowe. 2020. Osprey: A mmWave Approach to Tire Wear Sensing. In *Proc. of the 18th ACM MobiSys*. 28–41.
- [63] Raspberry Pi Foundation. 2021. Teach, Learn and Make with Raspberry Pi - Raspberry Pi. <https://www.raspberrypi.org/>. Accessed: 2021-03-02.
- [64] Robert Bosch Power Tools GmbH. 2017. D-tect 120 Wallscanner Professional. <https://www.bosch-professional.com/om/en/products/d-tect-120-wallscanner-0601081300>. Accessed: 2021-02-22.
- [65] Robert Bosch Tool Corporation. 2016. Bosch GLM 50 C. <https://www.boschtools.com/us/en/boschtools-ocs/laser-measures-glm-50-c-128522-p/>. Accessed: 2021-03-02.
- [66] Fabio Rocca, Ciro Cafforio, and Claudio Prati. 1989. Synthetic Aperture Radar: A New Application for Wave Equation Techniques. *Geophysical Prospecting* 37, 7 (1989), 809–830.
- [67] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 234–241.
- [68] Mohamed A. Safan, Zeinab A. Etman, and Alaa Konswa. 2019. Evaluation of Polyurethane Resin Injection for Concrete Leak Repair. *Case Studies in Construction Materials* 11 (2019), 1–12.
- [69] Florence Sagnard, Faroudja Bentabet, and Christophe Vignat. 2005. In Situ Measurements of the Complex Permittivity of Materials using Reflection Ellipsometry in the Microwave Band: Experiments (Part II). *IEEE Transactions on Instrumentation and Measurement* 54, 3 (2005), 1274–1282.
- [70] Screening Eagle Technologies. 2020. Proceq GPR Live - Ground Penetrating Radar. <https://www.proceq.com/compare/proceq-gpr-live/>. Accessed: 2021-02-22.
- [71] Longfei Shanguan and Kyle Jamieson. 2016. Leveraging Electromagnetic Polarization in a Two-Antenna Whiteboard in the Air. In *Proc. of the 12th CoNEXT*. 443–456.
- [72] S. Shihab and W. Al-Nuaimy. 2005. Radius Estimation for Cylindrical Objects Detected by Ground Penetrating Radar. *Subsurface Sensing Technologies and Applications* 6, 2 (2005), 151–166.
- [73] Mercedes Solla, Susana Lagüela, Norberto Fernández, and Iván Garrido. 2019. Assessing Rebar Corrosion through the Combination of Nondestructive GPR and IRT Methodologies. *Remote Sensing* 11, 14 (2019), 1705.
- [74] Gangbing Song, Haichang Gu, and Yi-Lung Mo. 2008. Smart Aggregates: Multi-Functional Sensors for Concrete Structures—A Tutorial and a Review. *Smart materials and structures* 17, 3 (2008), 033001.
- [75] Georger W. Stimson. 1998. Airborne Radar. *New Jersey: SciTech* (1998), 317–322.
- [76] Karthick Subbiah, Saraswathy Velu, Seung-Jun Kwon, Han-Seung Lee, Natarajan Rethinam, and Dong-Jin Park. 2018. A Novel In-Situ Corrosion Monitoring Electrode for Reinforced Concrete Structures. *Electrochimica Acta* 259 (2018), 1129–1144.
- [77] Richard K. Sutton. 2014. Aesthetics for Green Roofs and Green Walls. *Journal of Living Architecture* (2014), 1–19.
- [78] Vayyar Imaging. 2020. Walabot. <https://walabot.com/>. Accessed: 2021-06-17.
- [79] Ju Wang, Liqiong Chang, Shourya Aggarwal, Omid Abari, and Srinivasan Keshav. 2020. Soil Moisture Sensing with Commodity RFID Systems. In *Proc. of the 18th ACM MobiSys*. 273–285.
- [80] Yichen Wu, Yair Rivenson, Hongda Wang, Yilin Luo, Eyal Ben-David, Laurent A Bentolila, Christian Pritz, and Aydogan Ozcan. 2019. Three-Dimensional Virtual Refocusing of Fluorescence Microscopy Images Using Deep Learning. *Nature Methods* 16, 12 (2019), 1323–1331.
- [81] Lei Yang, Qiongzhen Lin, Xiangyang Li, Tianci Liu, and Yunhao Liu. 2015. See through Walls with COTS RFID System!. In *Proc. of the 21st ACM MobiCom*. 487–499.
- [82] Shuocho Yao, Yiran Zhao, Aston Zhang, Lu Su, and Tarek Abdelzaher. 2017. DeepIoT: Compressing Deep Neural Network Structures for Sensing Systems with a Compressor-Critic Framework. In *Proc. of the 15th ACM SenSys*. 1–14.
- [83] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. Gradient Surgery for Multi-Task Learning. In *Proc. of the 34th NIPS*. 1–13.
- [84] Chi Zhang, Feng Li, Jun Luo, and Ying He. 2014. iLocScan: Harnessing Multipath for Simultaneous Indoor Source Localization and Space Scanning. In *Proc. of the 12th ACM SenSys*. 91–104.
- [85] Diana Zhang, Jingxian Wang, Junsu Jang, Junbo Zhang, and Swarn Kumar. 2019. On the Feasibility of Wi-Fi Based Material Sensing. In *Proc. of the 25th ACM MobiCom*. 1–16.
- [86] Jianren Zhang, Hui Peng, and C.S. Cai. 2013. Destructive Testing of a Decommissioned Reinforced Concrete Bridge. *Journal of Bridge Engineering* 18, 6 (2013), 564–569.
- [87] Wei Zhang, Qiaofeng Zheng, Ashraf Ashour, and Baoguo Han. 2020. Self-Healing Cement Concrete Composites for Resilient Infrastructures: A Review. *Composites Part B: Engineering* 189 (2020), 107892.
- [88] Tianyue Zheng, Zhe Chen, Chao Cai, Jun Luo, and Xu Zhang. 2020. V²iFi: in-Vehicle Vital Sign Monitoring via Compact RF Sensing. In *Proc. of the 22nd ACM UbiComp*. 70:1–27.
- [89] Tianyue Zheng, Zhe Chen, Shuya Ding, and Jun Luo. 2021. Enhancing RF Sensing with Deep Learning: A Layered Approach. *IEEE Communications Magazine* 59, 2 (2021), 70–76.
- [90] Xiaoxiang Zhu, Sina Montazeri, Mohsin Ali, Yuansheng Hua, Yuanyuan Wang, Lichao Mou, Yilei Shi, Feng Xu, and Richard Bamler. 2021. Deep Learning Meets SAR: Concepts, Models, Pitfalls, and Perspectives. *IEEE Geoscience and Remote Sensing Magazine* (2021).
- [91] Yanzi Zhu, Yibo Zhu, Ben Y. Zhao, and Haitao Zheng. 2015. Reusing 60GHz Radios for Mobile Radar Imaging. In *Proc. of the 21st ACM MobiCom*. 103–116.