See No Evil, Hear No Evil, Feel No Evil, Print No Evil? Malicious Fill Patterns Detection in Additive Manufacturing

Christian Bayens, *Georgia Institute of Technology;* Tuan Le and Luis Garcia, *Rutgers University;* Raheem Beyah, *Georgia Institute of Technology;* Mehdi Javanmard and Saman Zonouz, *Rutgers University*

https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/bayens

This paper is included in the Proceedings of the 26th USENIX Security Symposium

August 16–18, 2017 • Vancouver, BC, Canada

ISBN 978-1-931971-40-9

Open access to the Proceedings of the 26th USENIX Security Symposium is sponsored by USENIX

See No Evil, Hear No Evil, Feel No Evil, Print No Evil? Malicious Fill Pattern Detection in Additive Manufacturing

Christian Bayens
Georgia Institute of Technology

Tuan Le
Rutgers University

Luis Garcia
Rutgers University

Raheem Beyah
Georgia Institute of Technology

Mehdi Javanmard
Rutgers University

Saman Zonouz
Rutgers University

Abstract

Additive Manufacturing is an increasingly integral part of industrial manufacturing. Safety-critical products, such as medical prostheses and parts for aerospace and automotive industries are being printed by additive manufacturing methods with no standard means of verification. In this paper, we develop a scheme of verification and intrusion detection that is independent of the printer firmware and controller PC. The scheme incorporates analyses of the acoustic signature of a manufacturing process, real-time tracking of machine components, and post production materials analysis. Not only will these methods allow the end user to verify the accuracy of printed models, but they will also save material costs by verifying the prints in real time and stopping the process in the event of a discrepancy. We evaluate our methods using three different types of 3D printers and one CNC machine and find them to be 100% accurate when detecting erroneous prints in real time. We also present a use case in which an erroneous print of a tibial knee prosthesis is identified.

1 Introduction

Additive Manufacturing (AM), also known as 3D printing, is an emerging field that shows promise in reducing waste, time, and infrastructure needed in a manufacturing process. Many major companies including Ford, GE, Airbus, SpaceX, Koenigsegg, and NASA are currently utilizing AM for both prototyping and production-quality manufacturing [43, 2, 1, 25, 15, 24]. Additionally, AM has been employed as a useful tool for printing medical implants [9], and cutting edge research is underway on producing food, drugs, and living tissue using AM techniques [4, 21]. Across industries, AM is expected to reach a market potential of 50% by 2038 [53].

Because of this potential for wide-spread use of AM in the coming decades, work has begun on understanding the security challenges that are unique compared to traditional manufacturing and cyber-physical security. Mark Yampolskiy, *et al.* [55] outlined a taxonomy for the potential of the misuse of a 3D printer as a weapon (3D-PaaW). In their paper, they identify the elements which may compromise or manipulate an AM environment, the targets of attack (printed object, printers, or environment), and the parameters for understanding the potential effectiveness of a given attack.

In this paper, we focus on the use of a 3D-PaaW to manipulate the physical properties of a printed object through manipulation of the object specifications, manufacturing parameters, and/or source material. According to the taxonomy described by Yampolskiy, *et al.* each of these are classified as attacks which would be achievable by an adversary through the manipulation of printer firmware or the controller PC. It has been shown that structural integrity can easily be compromised by introducing slight modifications in the model, e.g., a minuscule void injected into a manufactured dog bone can reduce the yield load by 14 percent [48].

In order to combat these forms of attack, we propose three methods of verification of design parameters that utilize analysis of the acoustic signal, embedded materials, and spatial position of machine components. These are chosen because they provide information about the manufactured design *without* access to the STL file or the G-code instructions[1] read by the printer. We do not consider our techniques to be a panacea for all verification needs. They are meant to be complementary to domain-specific verification methods. In some cases, this may be means of saving costs, e.g., by detecting malicious prints in real-time and ending them at the onset of a detection. In other cases, this may be a means of ensuring safety, e.g., by detecting malicious materials or designs before

[1] An STL file is a STereoLithography file for CAD software used in 3D printing. G-code is the set of actual instructions for 3D printers that are generated for particular models given an STL file and the print configuration, e.g., print speed and infill density.

the print is used. Throughout the course of this paper, we will consider the use case of printing the tibial portion of a knee prosthesis.

Our contributions are as follows:

- A multi-layered approach to the verification of design specifications, manufacturing parameters, and materials used in an AM.
- Proposed implementations of aforementioned approach for in-house and third-party AM producers.
- A case study of a scenario in which a malicious print of a medical prosthetic is identified.

The paper is organized as follows. We first provide a background in AM verification along with a system overview and threat model in section 2. We then provide details for the different types of verification methods that we proposed in section 3. In section 4, we evaluate the effectiveness of the combined verification scheme on a malicious print of a tibial knee implant. In section 5 we discuss the implementation and limitations of the verification scheme. We conclude in section 6 and discuss future work.

2 Background and System Model

In this section we discuss the previous efforts related to side-channel analysis of AM and verification of the physical models. We then provide a system overview of our approach as well as the threat model that will be used for the rest of the paper.

2.1 Side-Channel Analysis

In this paper we provide a means of verification by utilizing the various side-channels of the printing process. We also use materials science based verification to verify that the intended physical model is printed. As such, we first review previous efforts that have been made for the analysis of the side-channels involved in the AM process. We then provide a brief review on materials-based verification techniques like Raman spectroscopy and computed tomography (CT).

Acoustic, Magnetic, and Motion Sensing. KCAD [11] provided the first method of using the analog emissions of AM processes for the purpose of detecting so-called zero-day kinetic cyber-attacks. However, the work utilizes only one 3D printer and only investigates attacks in which simple variations in the exterior design. The paper also lacks any means of verifying the printed materials post-manufacturing. The focus of the majority of previous work on the analysis of side-channels from 3D printers used in AM has been its usefulness in obtaining intellectual property. Chen Song, *et al.* [44] and Avesta Hojjati, *et al.* [22] each showed that the array of sensors available on a modern smart phone can be leveraged to re-create designs produced from 3D printers or CNC machines. The sensors used in each study to collect side-channel data included the microphone, magnetometer, and accelerometer. Each group was able to reconstruct simple printed designs using supervised machine learning and manual analysis of sensor signals respectively. However, each group was only able to reconstruct very simple shapes such as two-dimensional outlines of airplanes or keys with no fill structure.

Beyond 3D printing and manufacturing, acoustic signals have also been shown to be useful in a growing number of security applications. As an example, Guri Mordechai, *et al.* [19] showed that information can be transmitted from a speakerless PC using information embedded in the sound of a cooling fan. Likewise, accelerometers have been used across industries as quality control sensors in CNC machines [31].

2.2 Physical Model Verification

The physical model that is printed from the AM machines are typically verified in a manner specific to the domain, such as mechanical strength testing [48]. Chien, *et al.* [12] use several techniques such as surface morphology characterization to verify 3D-printed tissue scaffolds. Furthermore, several solutions have been presented as preventative measures to future physical failures, such as the solution presented by Stava, *et al.* [45] for detecting and correcting models prior to being printed. However, these only correct the models that are being sent to the printer and do not verify the actual physical model in the event that the printer itself is compromised.

Imaging Analysis. We will now discuss the background for two modalities used for observing the composition of materials that will be explored in this paper for the verification of 3D printed models. It is important to note that we do not consider these modalities to be the most effective imaging techniques nor the most cost-effective solutions. As we will discuss in section 4, we chose these two modalities as they were readily available and are generalizable. Both solutions will act as a template for imaging techniques that are used to identify embedded materials. The choices for both the imaging technique and the associated embedded materials will be specific to the context in which they are applied.

Raman Spectroscopy. Surface-enhanced Raman spectroscopy (SERS) has been shown to be sensitive to single-molecule detection [35, 28, 34, 30]. Nie, *et al.* [35] have shown that silver colloidal nanoparticles can be used to amplify the spectroscopic signature of ad-

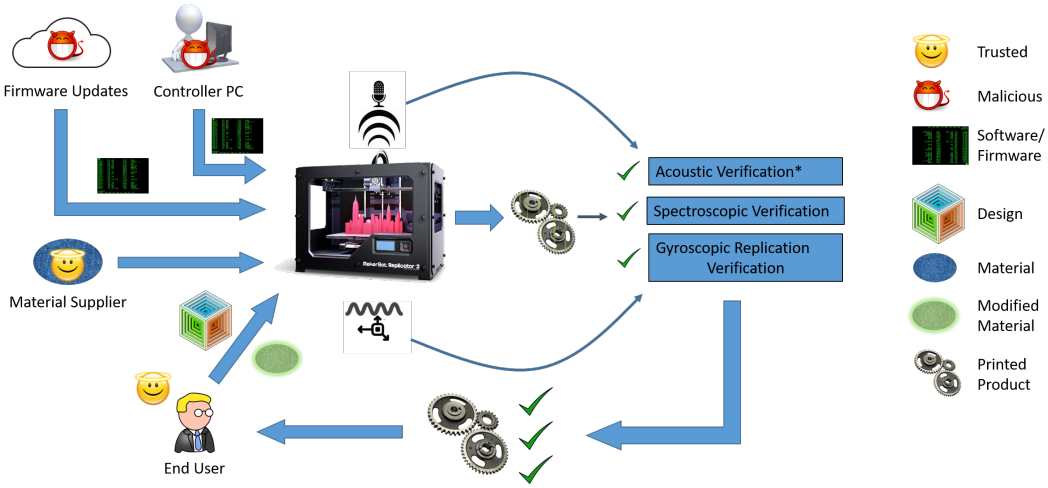

Figure 1: System Model.

sorbed Rhodamine 6G (R6G) and enable the single R6G molecule detection at room temperature. Furthermore, the sizes and shapes of the colloids enhance the spectral responses at different plasmon bands [36, 37]. We find that this technique can be utilized for post-production verification of 3D printed objects. By embedding a series of detectable markers of contrast agents in SERS at specific location within the 3D printed object, the SERS process would be able to reconstruct the model and verify the integrity of the internal structure of an object.

Computed Tomography. CT is typically used in medical applications to enable doctors to view precise images of their patients' internal organs [26]. Additionally, CT scanning also has been used in a wide variety of applications for verifying structural integrity. Cnudde, *et al.* [13] discuss the application of CT scanning in the context of geomaterials. Akin, *et al.* [5] also discuss the use of CT as a non-destructive method for imaging multiphase flow in porous media in the context of petroleum engineering research. Similarly, Alymore [7] discusses how CT scanning was used as a non-destructive method for studying soil behavior and soil/plant/water relations in space and time. In this study, we utilize CT in a similar fashion to construct models and verify the integrity of completed objects.

2.3 System Model

Figure 1 provides an overview of the system model that includes all verification techniques presented in this paper. Our system assumes that there is an end user with a 3D model design. The design will be printed on a 3D printer that is controlled by a controller PC. The 3D printer may or may not be controlled by a third party entity. The end user will send her design to be printed. Throughout the printing process, the object will be verified using three verification layers. The first two layers are achieved through acoustic side-channel analysis and spatial sensing which analyze the sound and physical position of printing components respectively. The third layer is that of materials verification in which imaging techniques are used to verify that the print is made from the proper material and printed correctly.

The end user may supply her own modified set of materials to the printer so that physical model verification may be performed upon completion. The goal is to embed special materials into the filament that is used in 3D printing. The modified filament can be used for materials verification purposes.

For the remainder of the paper, acoustic side-channel verification, spatial side-channel verification and materials verification are referred to as the acoustic layer, spatial layer, and material layer respectively.

2.4 Threat Model

The threat model assumes that the attacker has full knowledge of both the printer and its control software. If a third party manufacturer or affiliate of the user is involved, they are trusted as an organization. Therefore, they are willing to provide information about the print for verification. However, malicious entities may include network intruders, disgruntled employees, or other insider threats. The attack is carried out such that the printer behaves maliciously despite being sent G-code [2] for a non-malicious print. Meanwhile, the controller PC indicates that the print is being carried out correctly. This attack is feasable using a a cyber-physical rootkit such as Harvey described by Garcia, *et al.* [18].

[2] G-code is the set of instructions interpreted by a 3D-Printer, CNC, or other machine that includes information about motion direction, speed, and other operations.

It is also assumed that training prints may be performed under supervised circumstances in which it may be reasonably assumed that no attack is taking place. This may be achieved by a direct connection between the controlling machine and the printer via USB. The materials supplier shown in Figure 1 is assumed to be trusted. Untrusted materials suppliers are beyond the scope of this paper. For the materials-based verification, the modified filaments with the embedded materials are to be supplied directly by the end user. Furthermore, all communication channels among trusted entities are assumed to be secure.

2.5 Use Case: Prosthetic Tibial Implant

For a specific use case example, the tibial implant portion of a prosthetic knee was chosen. Unlike the titanium alloy component of the prosthetic knee that attaches to the femur, the tibial portion of the implant is made from polyethylene and has been identified as a component that could easily be manufactured through AM [9, 3]. Furthermore, the knee undergoes more mechanical stress than any joint [42]. Thus much research has been conducted which describes the medical implications of its wear and tear [50, 27]. Therefore, an attack is considered in which alterations are made to the internal structure of tibial knee implant that would dramatically increase the rate of wear.

3 Verification Layers and Implementation

The main focus of this paper is to verify the unseen internal fill structure present in all 3D printed objects. When a print is converted from a design on a computer to G-code instructions for a 3D printer or CNC, an internal structure for the physical product must be generated. These can range from low density for prototyping or non-load bearing prints to high density for load bearing or industrial use. The fill itself may take on a honeycomb pattern, rectilinear pattern, or other various patterns as specified by the user. Failure to produce the proper internal fill will render a final product that may externally look like the design intends, but fails to provide other required physical characteristics.

In order to develop a robust verification scheme, methods were needed that would allow for real-time identification and visualization of potentially malicious prints as well as visualization of a completed print to ensure its usability. Analysis of the acoustic side-channel was chosen as a non-intrusive method of identification. Instead of using traditional machine learning methods as have been used before, we use an audio classification scheme similar to popular apps used for identifying music. For real-time visualization, a method of tracking the moving components of a printer or CNC machine was determined to be a useful way of understanding the process without relying on control software. Finally, methods were borrowed from materials science by which the internal structure of an already completed print may be observed in a non-destructive way.

3.1 Side-Channel Verification

The side-channel analysis verification layers provide a means of verifying printed models in real-time. The goal is to infer as much information as possible from the given side-channels, but we do not expect each modality to be able to verify the entire print in itself. We will first discuss the experimental setup for each side-channel modality.

Acoustic Layer. As a physical byproduct of nearly any mechanical process, acoustic signals have been explored as a method of understanding information being processed by both traditional printers [8] and 3D Printers used in AM [44, 22, 11]. Because traditional printing methods now rely on lasers or ink jets, the information obtained from these is minimal. However, 3D printers will continue to rely on various actuators and fans for the foreseeable future which produce useful acoustic data. This is especially true for large-scale implementations of the technology.

In this verification layer, we assume that a particular design with a given infill structure will be printed multiple times. We use an open source audio classifier similar to the Shazaam [6] or SoundHound Applications. Using a training audio file, it locates noise-resistant peak frequencies and their temporal location within the file. It then locates frequency peaks in the test data that match the location, frequency, and spacing from other peaks. When a test file is identified, it is accompanied by a confidence score among other information. The confidence score indicates the number of peaks that the test has in common with the training data.

For AM verification, we use a single print as a training set by recording it with a microphone to obtain an audio file. Because even a simple print can take many minutes, the resulting file is separated into a number of segments of a given length (some number of seconds) and indexed in ascending order. Each indexed segment of the print is then trained as a different "song" and stored in a database. In many machine learning schema, common practice is to train multiple sets of data. However, because acoustic classification involves one-to-one comparison of audio files, a single-file training set is appropriate.

Test data is collected using the same method as training data and split into segments of the same length. Each indexed segment is then classified independently and a

confidence score is returned. The confidence score represents the number of frequency peaks that a given file has in common with the training file. Verification that a repeated print is unaltered from the training set is determined in two ways:

1. The classification results are such that the index values appear in ascending order. If they are out of order, it is likely that a change has been made.

2. The confidence score of one or more indexed classification results falls below a given threshold value. The threshold value is referred to as the confidence threshold (CTh) for the remainder of the paper. Its value is optimized manually for each printer to maximize the true positive rate and minimize the false positive rate.

With this, a print will be considered verified if each indexed audio file is classified correctly, in the correct order, and with confidence values greater than the CTh. A non-verified print conversely will be classified but out of order or with one or more confidence values less than CTh.

To test this method, two designs, shown in Figure 2 are used throughout this paper. They are described as a Rectangular Prism (right) and a Top Hat (left). Each was printed several times with "Honeycomb" and "Rectilinear" fill patterns of 20%, 40%, and 60% density. For each print style, a single set of audio data was split and stored in a unique database as described above.

In order to derive quantitative results to the test classifications, we assign a "score" to each segment of the audio data which are defined as follows:

- If a segment is in proper sequence and the confidence value is greater than CTh, its score is equal to that of the confidence value.

- If a segment is out of sequence, its score is equal to $-1 * \text{confidence value}$.

- If a segment is in sequence, but the confidence value is less than CTh, its score is set equal to $-1 * \text{confidence value}$.

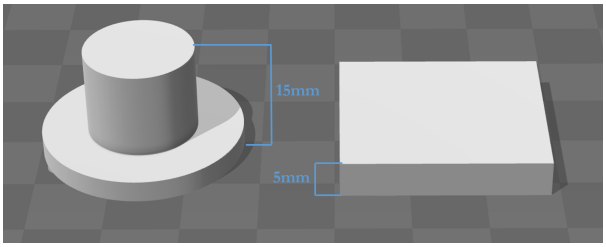

Figure 2: 3D Printed models described as (left) Top Hat and (right) Rectangular Prism.

If a negative score is calculated for any segment of the sliced audio file, a positive error classification may be determined. If no negative values are calculated, a negative error classification is determined.

Sample results are shown in Figure 3. The print is a Rectangular Prism with a 20% density Honeycomb fill pattern. The top chart shows the averaged results of three known negative error classifications (true negatives). Each bar represents a 90 second slice of the printing data, and CTh is set to 35. Likewise, the bottom chart represents various positive error classifications (true positives) caused by incorrect fill densities or patterns. Each type of error is printed four times and the results are averaged. For errors involving the Honeycomb fill pattern with erroneous densities, a positive error classification is achieved within 270s or the first 60% of the print. For the erroneous Rectilinear fill pattern, positive error classification is achieved within 180s or 40% of the print. In each case, the first 90s slice is always receives high scores due to the fact that the design always starts with a 100% density fill of the first three layers. This is standard in 3D printing to ensure that the exterior is solid.

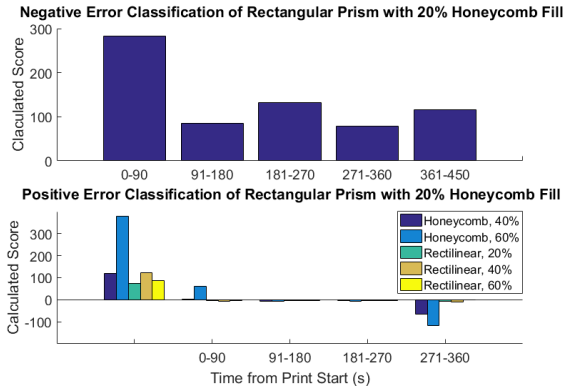

Figure 3: Classification example.

Spatial Sensing Layer. When performing 3D prints, it was found that the software used to monitor print progress simply displayed the progress of the G-code instructions being sent to the printer. This is regardless of the actual actions of the printer. The goal in setting up a spatial sensing verification scheme was to physically monitor the position of the printing nozzle with respect to the printing base, in order to observe their actual positions throughout the printing process.

The first consideration was to use a ride-along accelerometer such as those described in section 2. However, due to the double integration from acceleration to position and the noisiness of the accelerometer data, visual representations of the printer's path became prohibitively difficult to obtain.

With this in mind, a scheme was developed in which the a gyroscopic sensor was paired with a linear poten-

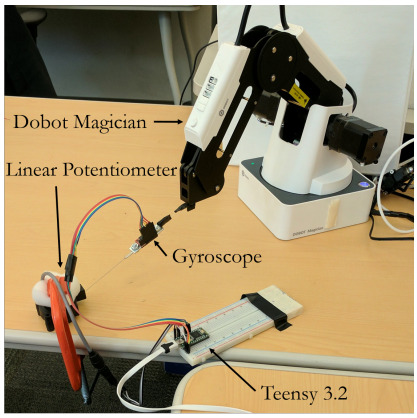

Figure 4: Spatial sensing setup with Unimeasure linear potentiometer model number LA-PA-10-N1N-NPC, SparkFun Triple Axis Accelerometer and Gyro Breakout, and Teensy 3.2 board.

tiometer in order to construct a set of spherical coordinates to describe the printer's motion. This proved more effective because no integration was needed for the data, and only simple moving average filtering was necessary to reduce noise.

To obtain these measurements, the following devices were used: a Unimeasure linear potentiometer model number LA-PA-10-N1N-NPC, a SparkFun Triple Axis Accelerometer and Gyro Breakout MPU-6050, and a Teensy 3.2 board. The experiments were conducted in a setup as shown in Figure 4 with a Dobot Magician desktop CNC and 3D Printer. For experimental purposes, the actual 3D printing extruder was removed and "dummy" prints were performed. The test prints were a single layer of a circular disk printed with Honeycomb and Rectangular fills each with a 20% and 40% density. Data is collected at a rate of 100Hz. In Figure 5, each print is shown as the G-code representation next to the reconstructed path of the printer. The data shown is smoothed using a moving average filter with a window of five.

3.2 Materials Verification

The objective of our materials-based verification is to embed contrast agents that will act as signature markers for particular prints without compromising the structural integrity of the original model. The contrast agents are chosen based on the materials as well as the scanning modalities. This approach is similar to the approach presented by Le, *et al.* [29] for privacy-preserving techniques for secure point-of-care medical diagnostics in which they used synthetic beads with different dielectric properties for user identification. In our case, we embed a single type of nanoparticle at different points in the printed model to generate a pattern specific to the model. This will allow us to ensure that the model was not modified by either an attacker who compromised the firmware and is duping the manufacturer, or a malicious insider who has physical access to the printing process. While it is arguable that embedded markers would change the integrity of the material itself, numerous studies have shown that the use of nanoparticles actually *improves* the materials' mechanical strength [54, 14, 17, 33].

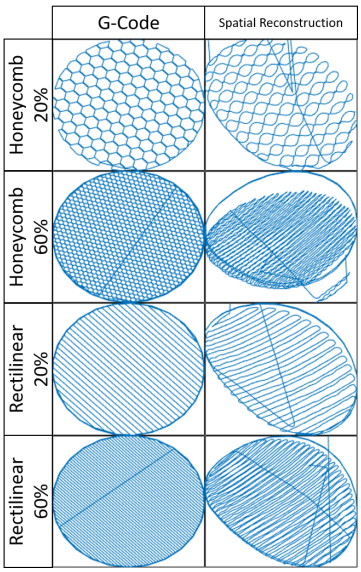

Figure 5: Comparison of G-code reconstruction to gyroscopic sensing reconstruction of single layers of various fill types and densities.

Here, we explore two types of scanning modalities: Raman spectroscopy and computed tomography (CT). Although both modalities are not necessarily cost-effective, our goal is to explore their effectiveness in our verification techniques. In both cases, we assume that the end user will provide the necessary materials to the manufacturer, who will be responsible for printing the model. The design sent to the manufacturer will not include any information about the embedded materials. We will now briefly discuss the different scanning modalities in detail.

Raman Spectroscopy. The first of the aforementioned modalities is Raman spectroscopy, which has been shown to be applicable for specific target identification and quantification [35, 28, 34, 30, 39, 47, 56]. The target sample is irradiated with a monochromatic light source such as laser. The majority of the scattering light has the same frequency of the incident light. This elastic scattering is called Rayleigh scattering. A small fraction of the scattering is inelastic. It has a small shift in photon frequency due to the energy transfer with the target molecules. When excited at a specific frequency, the target molecules can either increase or decrease in vibrational energy. Thus, the small fraction of the scattering light reduces (Stokes shift) or gains (anti-Stokes shift) equally the energy of the molecule vibration.

Due to to the unique covalent bonds and atomic mass of the each molecule, different molecules require specific excitation energy to change the molecule vibration [32]. The combination of multiple energy shifts creates the unique spectrum for each target molecule. The distinct spectra can be use to identify the target molecule in Raman spectroscopy.

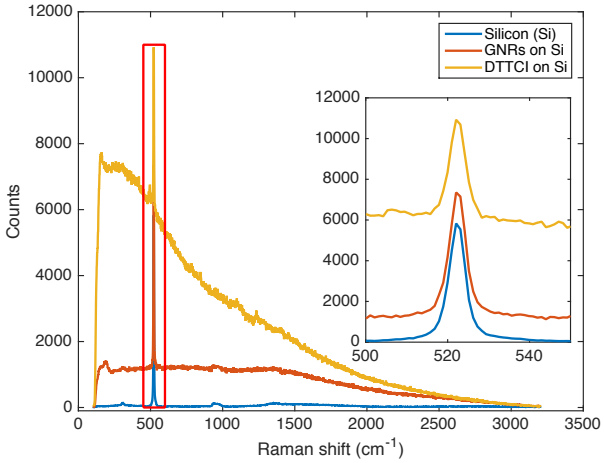

Figure 6: Raman scattering measurement of Silicon wafer with gold nanorods (GNRs) and 3,3'-Diethylthiatricarbocyanine iodide (DTTCI). The Raman spectrum of Si is amplified when using the enhancers.

Contrast agents in Surface enhanced Raman spectroscopy (SERS) can be used to amplify the Raman spectra of the target samples. As the electromagnetic wave (laser) irradiates the contrast agent molecules, it excites the localized surface plasmons on the rough surface. This results in the enhancement of electromagnetic fields near the surface [16, 10, 46]. The increase in intensity of the electromagnetic fields would also increase the intensity of Raman scattering. Thus, the Raman spectra is amplified. As a result, by coupling the contrast agents with the target molecules, SERS technique can be applied for identification of target molecules. Furthermore, SERS is also shown to be applicable for *in vivo* studying [40, 23]. Qian, *et al.* has shown that pegylated gold nanoparticles can be used to target tumor cells in live animals in an *in vivo* study.

In this study, we utilize gold nanorods (GNRs - *Sigma Aldrich*) and 3,3'-Diethylthiatricarbocyanine iodide (DTTCI - *Sigma Aldrich*) as the two different contrast agents in SERS detections to verify the material of the 3D printed object. The contrast agent can be embedded in the filament at specific locations for material identification. The internal structure of the 3D printed object can be verified using the embedded materials. Figure 6 shows the result of the standard Raman scattering measurement of the Silicon (Si) wafer and the Raman scattering of GNRs and DTTCI drop coat on top of the wafer. The Si wafer is used to calibrate the Raman instrument prior to the experiments. The Si Raman spectra has been studied thoroughly [38, 49, 41]. In Figure 6, the GNRs and the DTTCI amplified the signal response of the Si Raman scattering intensity.

Computed Tomography. The second scanning modality is a computed tomography (CT) scan. Just as in the SERS experiment, we needed to find an effective contrast agent that would allow us to view the embedded materials within the 3D printed model. Because it has been shown that gold works as an excellent contrast agent due to its X-ray density [20] and because we already had the materials at our disposal, we decided to reuse the GNRs as our contrast agent. Furthermore, the GNRs' biocompatability will allow us to apply our verification procedures to the tibial prosthesis.

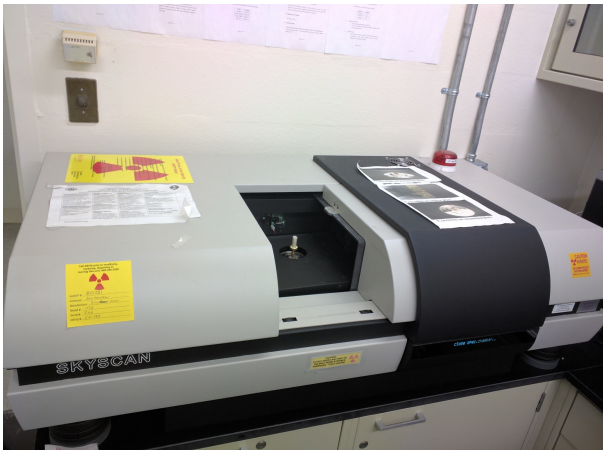

(a) Skyscan 1172 MicroCT scanner.

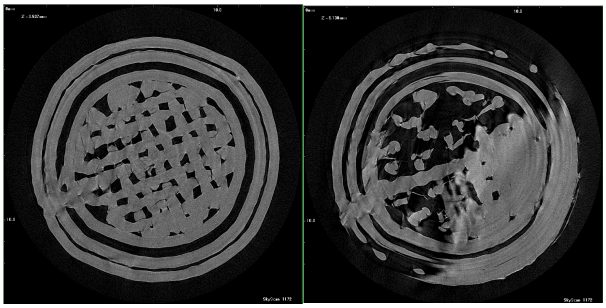

(b) ABS control print. (c) GNR layer print.

Figure 7: CT scan of ABS cylindrical tube with embedded GNRs.

We initially experimented with the use of GNRs as a contrast agent for CT scanning by embedding them in a simple 3D printed model. We developed and printed a cylindrical 3D model using a standard acrylonitrile butadiene styrene (ABS) filament as the control material of

the model. Multiple layers of ABS filament with embedded GNRs were deposited in between the bulk material.

Figure 7 shows the initial results of the 3D printed model with a layer of injected GNR filament. We performed a CT scan using a Skyscan 1172 MicroCT scanner. As the figure shows, the GNRs did indeed contrast with the ABS filament. This was sufficient to prove that GNRs could be used as a contrast agent for our printing use case. However, we will discuss in subsection 4.2 the limitations of the custom filament and as well as why we did not use the GNRs in our final evaluation.

4 Evaluation

In this section we evaluate the three-layered verification method. We describe the identification of a malicious print, the observation of the detected error, and the post-production materials verification. Then, we evaluate the effectiveness of the acoustic and spatial verification on the use case of a 3D printed tibial knee implant.

To quantify the accuracy of the results of the various tests, the data is fit into a logistic regression model with the binary dependent variable of "malicious print detected" or "no malicious print detected". From the model, we extract the probabilistic classification outcomes and create a receiver operating characteristic (ROC) curve. The area under the ROC curve (AUROC) is the metric used to predict classification accuracy.

Also, it is important to note that due to the fact that these machines are used to produce real 3D prints, large amounts of data were not practical to obtain. Furthermore, the imaging analysis techniques used for the materials verification were also time-consuming with limited availability. Therefore, sample sizes in this section will be significantly smaller than papers dealing with computer simulations.

4.1 Identification of Malicious Prints

In this section, we evaluate the usefulness of the proposed verification method in simply identifying an error in a potentially malicious print. This initial identification will be carried out primarily by the acoustic layer with redundancy in the spatial layer to reduce false classifications.

Classification Accuracy. In order to gain initial understanding of the parameters that affect the accuracy of the acoustic layer, several experiments were carried out with a small number of trials. The printers used in the tests were a Lulzbot Taz6, Lulzbot TazMini, and an Orion Delta. The AKG P170 condenser microphone was placed on a stand as close to the moving extruder head without being knocked over by the moving components of the printer. The audio classifier is called dejavu [52] and is an open-source project written in python.

In order to generate data useful for logistic regression, a vector of scores, **S**, is generated using the exact method as is described in subsection 3.1. For example, the components of **S** are what are shown in Figure 3. The vector **S** is of length n where $n = \lfloor \frac{\text{audio length}}{\text{audio slice length}} \rfloor$. We then calculate a print score, p, where

$$p = \sum_n S_n \, . \tag{1}$$

The value p associated with a given print now determines how likely the print is to be the same as the training print with higher values meaning more likely and lower values meaning less likely.

In Figure 8, the ROC curves are shown for the classification results of the Rectangular Prism design with Honeycomb and Rectilinear fills. The audio is segmented to 90 second and 120 second segments, each CTh = 35. The same original audio files are used whether the audio files are segmented to 90 seconds or 120 seconds. The Honeycomb and Rectilinear tests each consist of nine target prints and sixty malicious prints. The reason for the large number of known positive error classifications was that each print is considered an erroneous version of each other print.

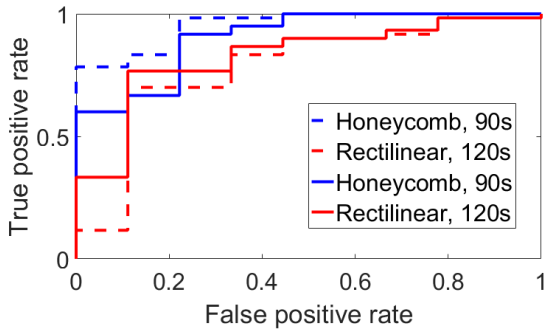

Figure 8: ROC Curve for Rectangular Prism, CTh = 35.

The poorest performance was an AUROC of 0.7815 for the rectilinear fill with the audio segmented at 90 seconds. That was determined to be unacceptable especially considering the high likelihood of false positives. To find an explanation for the poor classification, the G-code was inspected. Upon investigation of the G-code which was generated by Slic3r, it was found 9 lines which specified x and y coordinates along with the extrusion rate were repeated 12 times each out of 15 layers needed to complete the print in both the Rectilinear and Honeycomb fill patterns. Also, upon investigating sequentially repeated blocks of code, it was found that blocks of G-code describing three entire layers were repeated twice during

the course of the print. This symmetry was hypothesized to be the cause of the classification confusion.

To test this hypothesis, a second set of tests were conducted with the Top Hat design, which is asymmetrical along the z axis. The same number of prints was performed with Honeycomb and Rectilinear fill being sliced to 90s and 120s each and CTh set to 35. The ROC curve of these experiments are shown in Figure 9. Each sample consists of nine target prints and sixty malicious prints, and the same data is used for the 90 second audio slice length as the 120 second slice.

Upon investigation of the G-code, the only repeated lines were those that define the nozzle speed at the beginning and do not include extrusion. Furthermore, there are no blocks of G-code or layers that are entirely repeated verbatim. This is suspected to contribute greatly to the increased performance seen in Figure 9. Here, least AUROC is 0.9852 which is suitable for verification purposes. Between the 120 second and 90 second slice lengths, we see little change in performance. Although

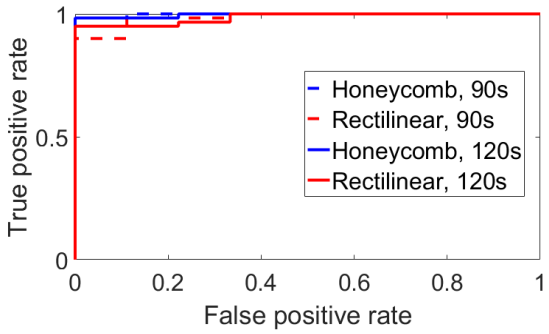

Figure 9: ROC Curves for Top Hat.

audio classification is shown here to be effective in identifying malicious prints, it is still susceptible to both false positives. By introducing data from the spatial layer, these may be reduced. For instance, Figure 10 compares the data from the x, y, and z axes of the 40% Honeycomb and 40% Rectilinear fills from Figure 5. Here, we see a significant difference between the two prints. Each frequency response has a similar shape, but the major features of the 40% Rectilinear fill are shifted to the right because the back-and-forth motion is not impeded by the creation of small Honeycomb structures.

For classification, the four most prominent peaks are used as features along with their locations. We conducted a test in which the target print was chosen to be the disk with 20% density Rectilinear fill shown above. All other prints were considered malicious. With this, we had 10 target prints and 12 malicious prints. Training using the linear regression model, an AUROC of 1.0 was achieved in differentiating between malicious and target prints.

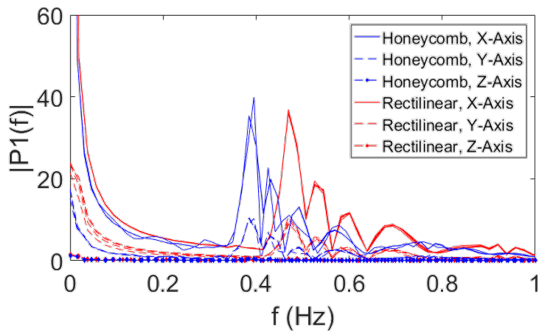

Figure 10: Comparison of the frequency response between a single layer of Honeycomb 40% fill and Rectilinear 40% fill. Four samples of each fill are compared.

While the spatial sensing layer is primarily for the purpose of print visualization, its role in conjunction with the acoustic layer allows for 100% accuracy in detecting malicious prints.

Varied Printer Models. In order to understand the effectiveness of audio classification for print verification on different printer models, several prints were performed on a Lulzbot TazMini and Orion Delta. Acoustic data recordings are obtained using the same microphone. In each print, a Top Hat design identical to the one described above was printed and the audio was sliced to 120s. The optimized CTh for the TazMini, Orion Delta, and Taz6 are 150, 20, and 35 respectively. The ROC curve results are shown in Figure 11. Because the Honeycomb and Rectilinear fill patterns are considered together, each data set consists of 18 target prints and 120 malicious prints. Consequently, the acoustic verification method is generalizable to printers of different sizes and configurations. The AUROC does not fall below 0.9542 in these tests.

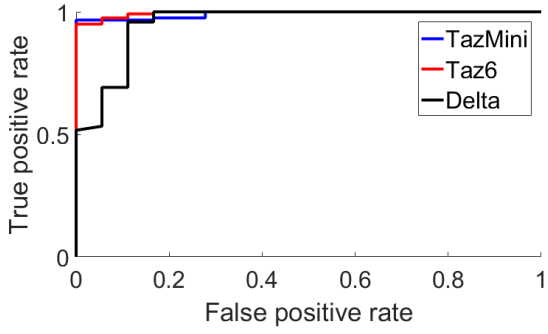

Figure 11: ROC curves for top hat design printed using a TazMini, Orion Delta, and Taz6 perint. Prints audio was sliced to 120 seconds and the confidence threshold is 150, 20, and 35 respectively.

Classification in Noisy Environments. Other experiments were conducted using an Afina H40 3D Printer with an eBoTrade Digital Voice Recorder wide-range microphone. This setup was in a noisy university makerspace with people talking near the printer. In this experiment, the classification accuracy suffered greatly (AUROC ≈ 0.5). Because it is shown that acoustic verification is useful on different types of printers above, we assume that the loss of classification accuracy is due to the noise in the environment. Also, because the microphone was wide range and not directional, the talking near the printer can be clearly heard. Therefore, in the implementation of this verification scheme it is important to use a directional microphone and noise isolation as much as possible.

4.2 Visualization of Malicious Prints

When a potentially malicious print is identified as described above, it is important to have the capability to visualize the potential threat. This visualization must be independent of the intended G-code which may be interpreted differently by malicious firmware. This is achieved in real time through use of the spatial sensing layer and in post-production by the materials inspection layer.

Real-Time Visualization. In the event that a potential malicious print is identified, a user has the capability of viewing the real-time print in progress through the spatial sensing as seen in Figure 5. By viewing the layer in progress, significant fill pattern changes such as those between the 20% Honeycomb and 20% Rectilinear fill are obvious. However, less obvious changes made to the print such as those between the 40% Honeycomb and Rectilinear fills are identifiable through FFT Analysis as in Figure 10. This is particularly true, as will be shown in subsection 4.3, if the user has access to the frequency response of a reference print.

While the spatial sensing layer is useful for identifying the type of fill pattern that is being maliciously generated, it is less useful for identifying if the design itself has been altered due to the warping that occurs in the data. This, however, is an easy issue to solve through the use of a webcam which can easily identify the shape of the design. In this sense, it may seem that spatial sensing may be replaced altogether by a webcam, but it is important that the latter uses far more data and does not readily provide information about the frequency response.

Post Production Visualization. The aforementioned materials-based verification methods are meant to be generalized for any scanning method that can detect the embedded contrast material within a 3D model. In our case, we chose Raman spectroscopy and computed tomography because those modalities were readily available to us at the time of evaluation.

Given the results shown in Figure 6, we concluded that the GNRs and DTTCI can be combined for use as a contrast agent in Raman spectroscopy. The contrast agents amplify the photon count across the Silicon spectrum in Raman spectroscopy. To echo the results shown in Figure 6 for the 3D printed disk, we use 10 nm diameter GNRs 780 nm absorption, and DTTCI 765 nm absorption *(Sigma Aldrich)* diluted in ethanol as the two distinct contrast agents. Each contrast agent is drop coated on the surface of the 3D printed disk. The Raman spectra of the blank 3D printed disk is also taken as the control data.

To emulate the filament with the embedded contrast agent, we produced the filament from ABS pellets using the filament maker *(Filabot)*. For the GNRs embedded filament, the ABS pellets are submerged in a GNR solution and left to dry. In this test, a 4 mL GNR solution was mixed with 12 g of pellets. Based on the information from the manufacturer, we naively calculated the number of GNRs per mL of solution to be approximately 7.284e11. Per 12 g of pellets, we can produce approximately 2 m of filament with a 2.5 mm diameter. The 3D printed disk has 50 µm in layer thickness. Therefore, for the area of 1 μm^2 on each layer of the 3D printed disk, there are approximately 4 GNRs particles. This approximation only serves as the estimation of the GNRs within the measurement area. Due to the non-uniform mixing of the the GNRs in the pellets, the distribution of GNRs within the 3D printed disk varies considerably. For the DTTCI embedded filament, while the quantity of DTTCI in the filament is not estimated, larger quantities of the DTTCI enhancer were available to produce the modified filament. The blank ABS filament is extruded using only ABS pellets.

Precise Embedding of Contrast Agents. In an ideal case, we would have the ability to embed the contrast agents or markers at precise Cartesian coordinates within the 3D printed models. However, for our proof of concept, we chose to simply create an ABS filament that was saturated in the GNRs or DTTCI throughout the entire spool of filament. The precise embedding of markers location is beyond the scope of current work. It can be explored in the near future. We then used a Lulzbot Taz dual extruder tool head to provide the capability of localize the embedded filament at precise locations.

In the following subsection, we evaluate the Raman spectra of the blank 3D printed disk, the 3D disk with GNRs or DTTCI drop coat on the surface, and the 3D printed disk with GNRs or DTTCI embedded filament. We wrote a simple C++ program that allowed the user to embed filament at desired locations by modifying the G-code where necessary, i.e., switching between the extruder nozzle containing the normal filament and the nozzle containing the GNR filament. The user can spec-

ify the beginning and end points of embedded material within the normal print path. This method was used for both the initial CT scan results as well as the final evaluation.

Imaging Analysis. In the evaluation using Raman spectroscopy, the 3D printed disk is excited with with 785 nm infrared light for 20 s per accumulation of data at 100 % power setting in Renishaw InVia micro-Raman system. Figure 12 shows the mean measurement results all data spectra of the 3D printed disks. Similar to the results from Figure 6, the spectrum of the 3D printed disk with DTTCI coated surface has significant improvement of photons counts across the spectrum comparing to the control data of the blank 3D printed disk. The spectra of the 3D printed disk from DTTCI embedded filament also shows the elevation of photons counts comparing to the control data. These spectra fall in between the spectra of the control data and the surface coated 3D printed disk. This conforms with the fact that the surface coated would accumulate more contrast agent at the measurement site comparing to the embedded filament. While the Raman spectroscopy can be used to quantify the concentration of the target particles, the elevation of the photons count in Figure 12 does not reflect the approximate distribution of contrast agent embedded in the filament. The measurement site in Raman spectroscopy might be a cluster or spare of contrast agent or markers. As mentioned above, the markers might not be uniformly distributed in the filament. This is confirmed in Figure 7c as a result of the MicroCT scanner. The high reflection in the CT scan shows the large cluster of the GNRs in the embedded filament. Due to the low resolution of the MicroCT scanner, the scan would not highlight the areas where the GNRs are sparsely distributed. While the Raman spectroscopy results of the GNRs embedded filament are not shown, the similar response can be discerned.

In classification of 3D printed blank ABS, GNRs embedded, and DTTCI embedded disk, mean and standard deviation of the spectra are used to distinguish the cluster of data set. Figure 12 shows the mean of the typical response of Raman spectra of 3D printed disk with blank ABS, DTTCI coated disk, and DTTCI embedded ABS filament. By observation, the greatest change of Raman shift is in the range of $100cm^{-1}$ and $800cm^{-1}$. The details of the Raman scattering separation can be seen in Figure 20 in Appendix A. This is in the range of $791.21nm$ and $837.60nm$ scattering; whereas the sample is irradiated at $785nm$. Therefore, this is the reasonable range of interest for Raman scattering for all data selection. By training the logistic regression model, the classification using mean and standard deviation shows 100 % accuracy against the blank ABS (226 samples) filament for both GNRs (179 samples) and DTTCI (71 samples) embedded filaments.

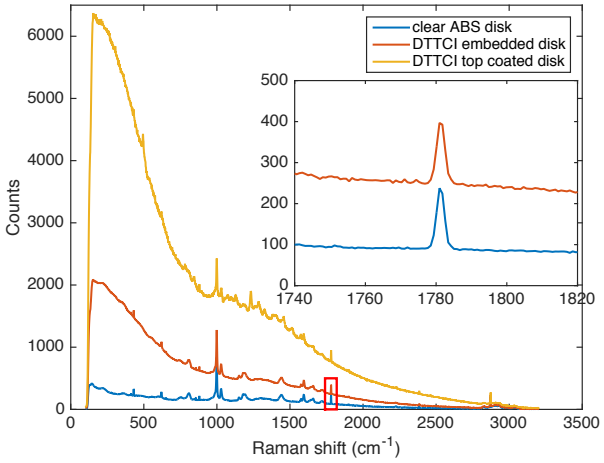

Figure 12: Mean measurement of Raman scattering of 3D printed disks using acrylonitrile butadiene styrene (ABS) filament and ABS with gold naonorods (GNRs) and 3,3'-Diethylthiatricarbocyanine iodide (DTTCI) embedded.

In Raman spectroscopy, the maximum setting depth penetration for the Renishaw InVia micro-Raman system is approximately 300 µm, we cannot verify the 3D printed object where the GNRs or DTTCI embedded filament is implanted further inside the object. Therefor, the Raman spectroscopy would not be sufficient for the verification that require depth. In further analysis, we use the MicroCT scanner to evaluate the internal structure of 3D printed objects.

The initial results for the CT scan approach presented in Figure 7 showed that although the GNRs embedded filament contrasted well in the CT scan, we could not rely on the custom filament due to the sparse distribution of the GNRs. We did not have the equipment nor the expertise to manufacture a heavily saturated filament.

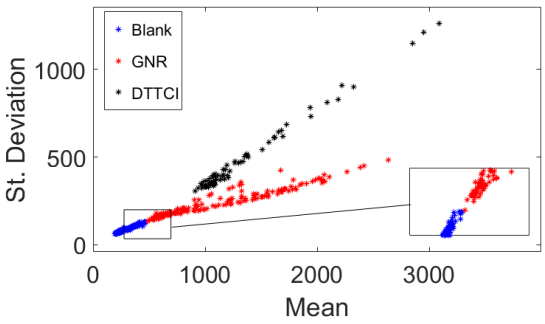

Figure 13: Classification of blank acrylonitrile butadiene styrene (ABS), gold nanorods (GNRs), and 3,3'-Diethylthiatricarbocyanine iodide (DTTCI) dye embedded filament in 3D printed disks.

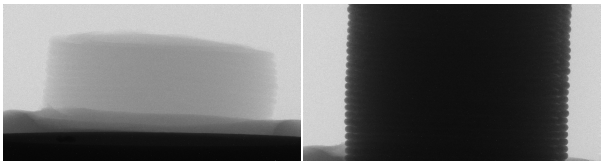

(a) PLA filament. (b) Stainless steel filament.

Figure 14: Comparison of X-ray densities of PLA and stainless steel filaments.

For a more precise proof of concept, we used commercially available stainless steel filaments where the filament is heavily saturated with stainless steel particles. Under the CT scanning, the steel particles would produce similar response to the GNRs due to high X-ray density. Although stainless steel is not biocompatible, it will serve as a substitute for the GNRs in order to provide precise visibility in the CT scan. Furthermore, we changed the control filament from ABS to polylactic acid (PLA) after comparing the densities in the CT scan. The X-ray properties of PLA versus ABS have been studied [51], but we confirmed our assumption after simple trial and error. Figure 14 highlights the contrast in X-ray densities between the PLA filament and the stainless steel filament. We will discuss in the subsequent section how we evaluated this approach on a tibial prosthesis.

4.3 Case Study: Prosthetic Knee

As described in subsection 2.5, a model of the tibial component of a prosthetic knee implant was used as a design for a use case test. Prosthetics differ slightly between patients, so we assume that malicious print identification is performed periodically with a known standard prosthetic design. Real-Time and post-production visualization are still performed on each print.

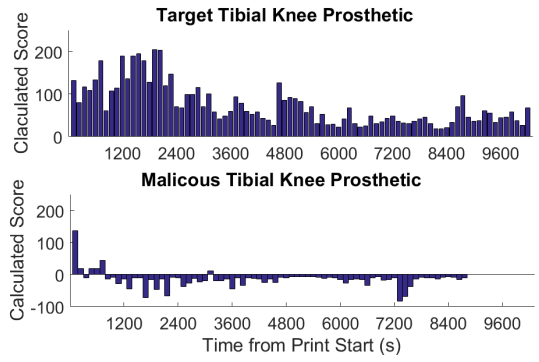

Figure 15: Comparison of target 60% Rectilinear Fill Tibial Prosthetic print acoustic classification (Top) vs. malicious 20% Honeycomb Fill (bottom). CTh = 0.

Error Identification. The acoustic verification results are shown in Figure 15 which shows the confidence values of both the target print and the malicious print. These results are gathered using the same technique as those described in section 3 with audio slices of length 120s and CTh = 0. By setting CTh = 0, we see that a positive error classification can be made within the first 360s of the print or the first 4% of the total known print time by only observing out-of-sequence index classifications. The CTh may be set to anything less than 18 without causing a false positive. Overall, acoustic error detection itself saves over 2 hours of print time and prevents a potentially harmful print from being completed. A detailed table of the results shown here can be found in Appendix B.

In Figure 16, the FFT of a target print and a malicious print are compared to a training print. Similar to Figure 10, the malicious print shows a different frequency response near 0.2Hz as highlighted by the lower box. The upper box highlights the closeness of the peaks between the training and target prints and the difference between those and the malicious print. The full print of the object requires 111 layers, so it would take less 1% of the time of the total print to identify the erroneous pattern once it begins.

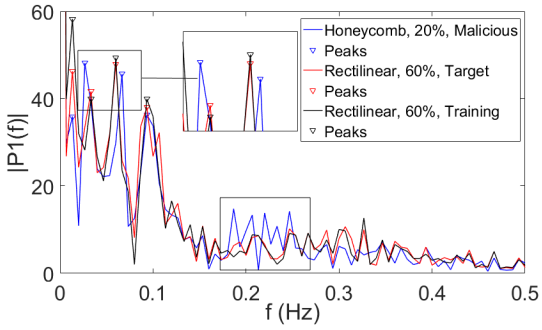

Figure 16: Comparison of x-axis frequency response for a layer of a layer of the tibial knee implant design.

Real-Time Visualization. In this test, the target print uses a 60% Rectilinear fill and the malicious print uses a 20% Honeycomb fill. In the attack, the visualization of the intended G-code remains unaltered for the user while the instructions sent to the printer are altered. The consequences of this attack would be to cause accelerated wear in the implant causing pain and financial loss for the victim who has the implant.

For the print identification and real-time visualization tests, a full sized prosthetic design is used. However, due to the size limitations of the MicroCT scanner, a significantly scaled down version of the same design is used.

The training, target, and attack prints were each recorded on the Lulzbot Taz6 printer. Due to the availability of the experimental setup, a single layer of each

of these prints was performed by the Dobot Magician for the visualization tests. The exact same G-code was used for the Dobot prints as in the Taz6 with the exception of the extruder being disabled and the speeds decreased to suit the capabilities of Dobot. It should be noted that spatial verification testing is entirely plausible on the Taz6 which has a moving base because the measurements describe the relative position between the nozzle and the base. This is regardless of whether that base is a stationary table or a moving part of the printer. It should also be noted that both acoustic and spatial verification would ideally be performed in tandem, but for testing purposes here, they are not.

Figure 17 shows the spatial verification visualization of, in order of left to right, a G-code visualization of the training print, a spatial reconstruction of the target print, and a spatial reconstruction of the malicious print. It is clear that the recreated target print uses a rectilinear fill at approximately the correct density while the malicious print differs significantly from the intended G-code. Due to the warping that occurs in the spatial reconstruction, a user would not be made aware if the shape of the print were altered by using this method alone.

Post Production Visualization. We only considered the CT scan approach for the post production visualization as the Raman spectroscopy would not be able to verify the internal structure of the tibial prosthesis due to its depth limitations. Figure 18 shows an X-ray scan of the front of a PLA tibial prosthesis with 2 infill layers of steel. Because we had to use a MicroCT scanner, the part of the tibial insert was scaled down to fit within a diameter of about 30 mm. The two large blotches of stainless steel are simple imperfections that mark points where the second extruder began printing.

Figure 19 compares the G-Code representation of the intended print of the top stainless steel layer–with the stainless steel path highlighted in red–versus the CT scan of that layer at a 15 μm/voxel resolution. The CT scan image is rotated about 45 degrees in comparison to the intended print. Furthermore, the small model had to be mounted on a bed of silicone polymer to hold it in place, so it is not completely level. Despite the imperfections of the printed model and the scans, it can be seen that the steel was properly embedded within the walls of the model and is clearly detectable against the PLA filament.

5 Discussion

In this section, we discuss the various methods of implementing the proposed verification scheme. We then briefly discuss its limitations.

Implementation. The three layer verification and malicious print detection scheme described here is most readily suited for a mass production AM scenario. In this setting, many different standard designs may be produced using the same equipment. If each design is printed identically, then the acoustic layer, spatial sensing layer, and materials verification layer may be applied to each individual print.

In a setting such as the one described for the case study in subsection 4.3, a base design may be modified for each print in order to adjust for biological parameters, etc. In this scenario, the user could train a known standard print and periodically test the printer for any malicious activity. This periodic test could include all three layers. Each specialized design, then, could be monitored using spatial and materials verification for real time and post production detection of malicious activity.

Finally, this verification scheme may be used in a scenario in which an end user sends a design to a third party to be printed. For the materials verification layer, she may send a specialized filament with embedded trackers to be used. If the object returns without the trackers or with trackers in the wrong locations, malicious activity may be detected. Also, using a secure live streaming connection, the user may receive data from the print in progress and perform any classification or analysis herself.

The experiments presented in this paper focus primarily on on the detection of subtle changes in the internal fill pattern. Therefore, it is logical that more significant changes such as holes in the fill pattern or changes in the overall design will be easily detected.

Limitations. As with any verification schema, the system proposed here is not without limitations. The immediately obvious limitation is that the ability to detect a deviation from a training print decreases as the similarity to the print increases. However, drawn to its logical conclusion, this means that an attacker wishing to exploit this limitation would be forced to change the design in such a small way as to not affect its usefulness. Another limitation could be the need for a training print. This may be a minor issue in the mass production scheme described above. In a scenario such as the production of prosthetics, however, the periodic checks for malicious activity may be seen as time consuming. Finally, if a third party printing service implements these methods, some cost overhead will incur from the purchase of microphones, sensors, etc. However, these costs are relatively cheap considering that any major equipment such as a spectroscope or CT scanner would be in the domain of the end user.

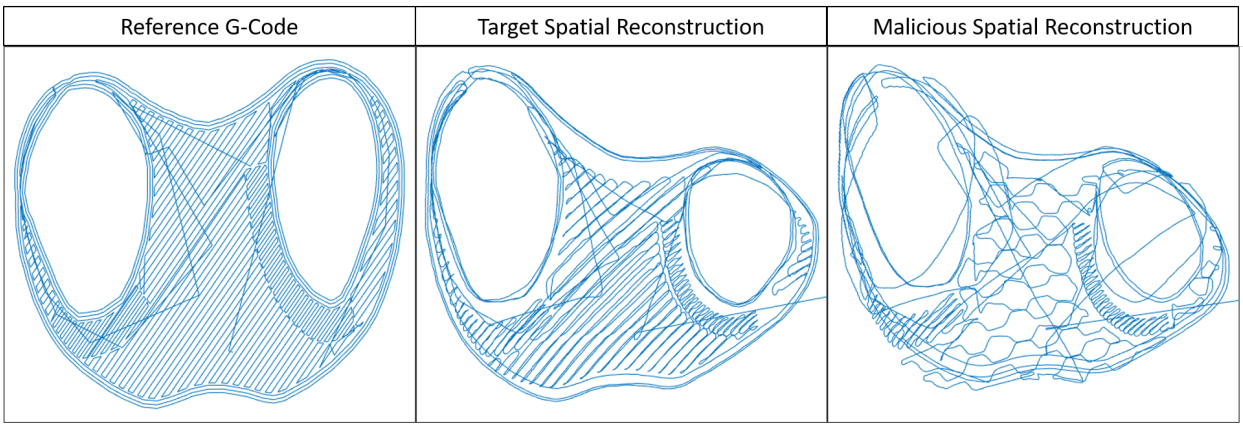

Figure 17: Comparison of target and malicious tibial knee implant prints. Left: G-code reconstruction of 60% Rectilinear fill, Middle: Spatial reconstruction of 60% Rectilinear fill, Right: Spatial reconstruction of malicious 20% Honeycomb fill.

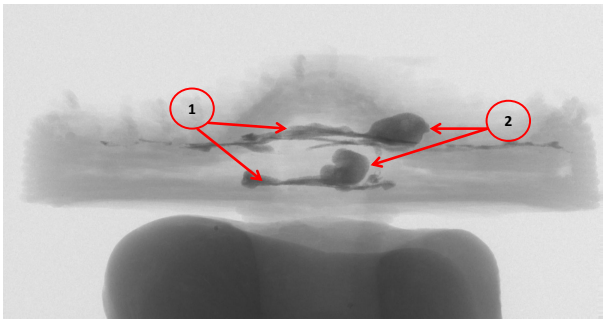

Figure 18: X-ray scan of front of PLA tibia with embedded stainless steel at a 15 μm/voxel size resolution. The first label shows the side view of the cross-sectional stainless steel infill, while the second label shows the two blotches where the stainless steel print began.

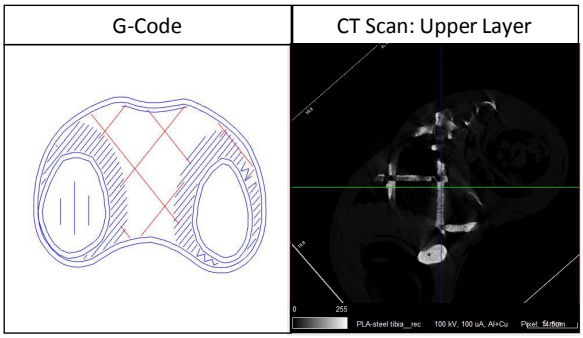

Figure 19: Comparison of G-code simulation of embedded steel (shown as red lines) versus CT scan of the printed model. The CT scan image is rotated about 45 degrees.

6 Conclusion

Three layers of verification for AM are presented for a case in which either a control PC or printer firmware is compromised. Acoustic verification uses audio classification to determine whether a print matches a previously known print. Spatial verification provides a visualization of the print in real time along with data for frequency analysis of the printing process. Materials verification determines whether the correct materials were used and whether indicator patterns appear in the proper locations. Each layer is independent of firmware or a controller PC.

Acoustic and spatial verification are found to be useful for confirming the intended fill pattern and density in a print, and material verification is found to be most useful in determining that the correct material is used and that the design is free of tampering.

Future work will include improving the acoustic and spatial classification methods so that they work independently of human interaction and in real-time. Similarly, the materials verification methods presented in this paper could be tuned for domain-specific solutions to be more precise. This would facilitate automated materials verification solutions.

Acknowledgement

We would like to thank the National Science Foundations (NSF - CNS 1453046) for their support of this work. Additionally, we would like to thank the following individuals for their technical assistance: Dr. Patricia Buckendahl at the Rutgers University MicroCT Imaging Facility, Erik Shuster at the Georgia Institute of Technology Invention Studio, as well as Sakshi Sadar of the Materials Science and Engineering Department at Rutgers University.

References

[1] Arconic strengthens 3d printing collaboration with airbus. *http://advancedmanufacturing.org/arconic-airbus-3d-printing-collaboration/*, Dec 2016.

[2] Hardware meets software in advanced manufacturing. *https://www.ge.com/stories/hardware-meets-software-advanced-manufacturing*, 2017.

[3] Knee replacement implant materials. *https://bonesmart.org/knee/knee-replacement-implant-materials/*, 2017.

[4] Natural machines: The makers of foodini - a 3d food printer making all types of fresh, nutritious foods. *http://www.naturalmachines.com/*, 2017.

[5] S Akin and AR Kovscek. Computed tomography in petroleum engineering research. *Geological Society, London, Special Publications*, 215(1):23–38, 2003.

[6] Avery Li-Chun Wang. An industrial strength audio search algorithm.

[7] LAG Aylmore. Use of computer-assisted tomography in studying water movement around plant roots. *Advances in Agronomy*, 49:1–54, 1993.

[8] Michael Backes, Markus Drmuth, Sebastian Gerling, Manfred Pinkal, and Caroline Sporleder. Acoustic side-channel attacks on printers. In *Proceedings of the 19th USENIX Conference on Security*, USENIX Security'10, pages 20–20. USENIX Association.

[9] Barry Berman. 3-d printing: The new industrial revolution. 55(2):155–162.

[10] Alan Campion and Patanjali Kambhampati. Surface-enhanced raman scattering. *Chemical Society Reviews*, 27(4):241–250, 1998.

[11] Sujit Rokka Chhetri, Arquimedes Canedo, and Mohammad Abdullah Al Faruque. Kcad: Kinetic cyber attack detection method for cyber-physical additive manufacturing systems. In *Proceedings of the 35th International Conference on Computer-Aided Design*, page 74. ACM, 2016.

[12] Karen B Chien, Emmanuella Makridakis, and Ramille N Shah. Three-dimensional printing of soy protein scaffolds for tissue regeneration. *Tissue Engineering Part C: Methods*, 19(6):417–426, 2012.

[13] Veerle Cnudde and Matthieu Nicolaas Boone. High-resolution x-ray computed tomography in geosciences: A review of the current technology and applications. *Earth-Science Reviews*, 123:1–17, 2013.

[14] Alfred J Crosby and Jong-Young Lee. Polymer nanocomposites: the nano effect on mechanical properties. *Polymer reviews*, 47(2):217–229, 2007.

[15] Alex Davies, 2014 Feb. 28, and 199 6. A swedish automaker is using 3d printing to make the world's fastest car.

[16] Martin Fleischmann, Patrick J Hendra, and A James McQuillan. Raman spectra of pyridine adsorbed at a silver electrode. *Chemical Physics Letters*, 26(2):163–166, 1974.

[17] Shao-Yun Fu, Xi-Qiao Feng, Bernd Lauke, and Yiu-Wing Mai. Effects of particle size, particle/matrix interface adhesion and particle loading on mechanical properties of particulate–polymer composites. *Composites Part B: Engineering*, 39(6):933–961, 2008.

[18] Luis Garcia, Ferdinand Brasser, Mehmet H. Cintuglu, Ahmad-Reza Sadeghi, Osama Mohammed, and Saman A. Zonouz. Hey, my malware knows physics! attacking plcs with physical model aware rootkit. In *24th Annual Network & Distributed System Security Symposium (NDSS)*, February 2017.

[19] Mordechai Guri, Yosef Solewicz, Andrey Daidakulov, and Yuval Elovici. Fansmitter: Acoustic data exfiltration from (speakerless) air-gapped computers.

[20] JF Hainfeld, DN Slatkin, TM Focella, and HM Smilowitz. Gold nanoparticles: a new x-ray contrast agent. *The British journal of radiology*, 2014.

[21] Jennifer Hicks. FDA approved 3d printed drug available in the US.

[22] Avesta Hojjati, Anku Adhikari, Katarina Struckmann, Edward Chou, Thi Ngoc Tho Nguyen, Kushagra Madan, Marianne S. Winslett, Carl A. Gunter, and William P. King. Leave your phone at the door: Side channels that reveal factory floor secrets. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 883–894. ACM.

[23] Xiaohua Huang, Ivan H El-Sayed, Wei Qian, and Mostafa A El-Sayed. Cancer cells assemble and align gold nanorods conjugated to antibodies to produce highly enhanced, sharp, and polarized surface raman spectra: a potential cancer diagnostic marker. *Nano letters*, 7(6):1591–1597, 2007.

[24] G. D. Janaki Ram, Y. Yang, and B. E. Stucker. Effect of process parameters on bond formation during ultrasonic consolidation of aluminum alloy 3003. 25(3):221–238.

[25] Foust Jeff. SpaceX unveils its 21st century spaceship.

[26] Avinash C Kak and Malcolm Slaney. *Principles of computerized tomographic imaging*. SIAM, 2001.

[27] D.J. Kilgus, J.R. Moreland, G.A.M. Finerman, T.T. Funahashi, and J.S. Tipton. Catastrophic wear of tibial polyethylene inserts. (273):223–231.

[28] Katrin Kneipp, Yang Wang, Harald Kneipp, Lev T Perelman, Irving Itzkan, Ramachandra R Dasari, and Michael S Feld. Single molecule detection using surface-enhanced raman scattering (sers). *Physical review letters*, 78(9):1667, 1997.

[29] Tuan Le, Gabriel Salles-Loustau, Laleh Najafizadeh, Mehdi Javanmard, and Saman Zonouz. Secure point-of-care medical diagnostics via trusted sensing and cyto-coded passwords. In *Dependable Systems and Networks (DSN), 2016 46th Annual IEEE/IFIP International Conference on*, pages 583–594. IEEE, 2016.

[30] Eric C Le Ru, Matthias Meyer, and Pablo G Etchegoin. Proof of single-molecule sensitivity in surface enhanced raman scattering (sers) by means of a two-analyte technique. *The journal of physical chemistry B*, 110(4):1944–1948, 2006.

[31] Richard L. Lemaster, Liya Lu, and Steve Jackson. The use of process monitoring techniques on a CNC wood router. part 2. use of a vibration accelerometer to monitor tool wear and workpiece quality. 50(9):59–64.

[32] Daimay Lin-Vien, Norman B Colthup, William G Fateley, and Jeanette G Grasselli. *The handbook of infrared and Raman characteristic frequencies of organic molecules*. Elsevier, 1991.

[33] Huinan Liu and Thomas J Webster. Mechanical properties of dispersed ceramic nanoparticles in polymer composites for orthopedic applications. *Int J Nanomedicine*, 5:299–313, 2010.

[34] Amy M Michaels, M Nirmal, and LE Brus. Surface enhanced raman spectroscopy of individual rhodamine 6g molecules on large ag nanocrystals. *Journal of the American Chemical Society*, 121(43):9932–9939, 1999.

[35] Shuming Nie and Steven R Emory. Probing single molecules and single nanoparticles by surface-enhanced raman scattering. *science*, 275(5303):1102–1106, 1997.

[36] Babak Nikoobakht and Mostafa A El-Sayed. Surface-enhanced raman scattering studies on aggregated gold nanorods. *The Journal of Physical Chemistry A*, 107(18):3372–3378, 2003.

[37] Christopher J Orendorff, Latha Gearheart, Nikhil R Jana, and Catherine J Murphy. Aspect ratio dependence on surface enhanced raman scattering using silver and gold nanorod substrates. *Physical Chemistry Chemical Physics*, 8(1):165–170, 2006.

[38] JH Parker Jr, DW Feldman, and M Ashkin. Raman scattering by silicon and germanium. *Physical Review*, 155(3):712, 1967.

[39] Dahu Qi and Andrew J Berger. Quantitative concentration measurements of creatinine dissolved in water and urine using raman spectroscopy and a liquid core optical fiber. *Journal of biomedical optics*, 10(3):031115–0311159, 2005.

[40] Ximei Qian, Xiang-Hong Peng, Dominic O Ansari, Qiqin Yin-Goen, Georgia Z Chen, Dong M Shin, Lily Yang, Andrew N Young, May D Wang, and Shuming Nie. In vivo tumor targeting and spectroscopic detection with surface-enhanced raman nanoparticle tags. *Nature biotechnology*, 26(1):83–90, 2008.

[41] H Richter, ZP Wang, and L Ley. The one phonon raman spectrum in microcrystalline silicon. *Solid State Communications*, 39(5):625–629, 1981.

[42] Cindy Schmidler. Knee joint anatomy, function and problems. *http://www.healthpages.org/anatomy-function/knee-joint-structure-function-problems/*, Dec 2016.

[43] Rick Smith. 8 hot 3d printing trends to watch in 2016.

[44] Chen Song, Feng Lin, Zhongjie Ba, Kui Ren, Chi Zhou, and Wenyao Xu. My smartphone knows what you print: Exploring smartphone-based side-channel attacks against 3d printers. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 895–907. ACM.

[45] Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomír Měch. Stress relief: improving structural strength of 3d printable objects. *ACM Transactions on Graphics (TOG)*, 31(4):48, 2012.

[46] Paul L Stiles, Jon A Dieringer, Nilam C Shah, and Richard P Van Duyne. Surface-enhanced raman spectroscopy. *Annu. Rev. Anal. Chem.*, 1:601–626, 2008.

[47] Clare J Strachan, Thomas Rades, Keith C Gordon, and Jukka Rantanen. Raman spectroscopy for quantitative analysis of pharmaceutical solids. *Journal of pharmacy and pharmacology*, 59(2):179–192, 2007.

[48] L Sturm, C Williams, J Camelio, J White, and R Parker. Cyber-physical vunerabilities in additive manufacturing systems. *Context*, 7(2014):8, 2014.

[49] Paul A Temple and CE Hathaway. Multiphonon raman spectrum of silicon. *Physical Review B*, 7(8):3685, 1973.

[50] Pieter-Jan T. K. Vandekerckhove, Matthew G. Teeter, Douglas D. R. Naudie, James L. Howard, Steven J. MacDonald, and Brent A. Lanting. the impact of coronal plane alignment on polyethylene wear and damage in total knee replacement: a retrieval study.

[51] GR Veneziani, EL Corrêa, MPA Potiens, and LL Campos. Attenuation coefficient determination of printed abs and pla samples in diagnostic radiology standard beams. In *Journal of Physics: Conference Series*, volume 733, page 012088. IOP Publishing, 2016.

[52] Drevo Will. Dejavu; available at `https://github.com/worldveil/dejavu`, 2017.

[53] Terry Wohlers. *Wohlers Report 2015: 3D printing and additive manufacturing state of the industry; annual worldwide progress report*. Wohlers Associates, 2015.

[54] Chun Lei Wu, Ming Qiu Zhang, Min Zhi Rong, and Klaus Friedrich. Tensile performance improvement of low nanoparticles filled-polypropylene composites. *Composites Science and Technology*, 62(10):1327–1340, 2002.

[55] Mark Yampolskiy, Anthony Skjellum, Michael Kretzschmar, Ruel A. Overfelt, Kenneth R. Sloan, and Alec Yasinsac. Using 3d printers as weapons. 14:58–71.

[56] Qingyuan Zhu, Robert G Quivey, and Andrew J Berger. Raman spectroscopic measurement of relative concentrations in mixtures of oral bacteria. *Applied spectroscopy*, 61(11):1233–1237, 2007.

APPENDIX

A Raman Spectroscopy Measurements

Figure 20 shows the Raman spectroscopy measurements of 3D printed disks of Raman scattering enhancers gold nanorods (GNRs), and Diethylthiatricarbocyanine iodide (DTTCI) embedded in acrylonitrile butadiene styrene (ABS) filament.

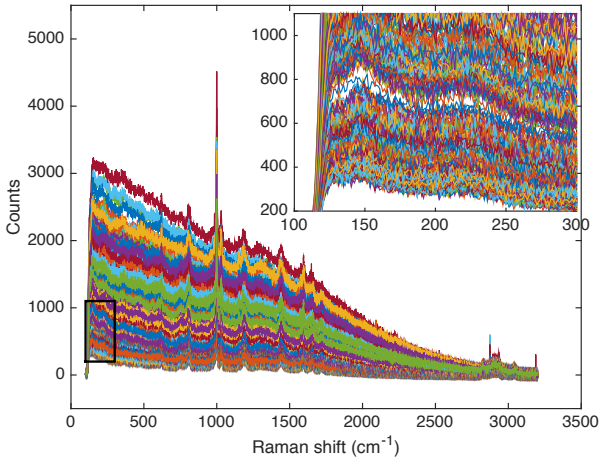

(a)

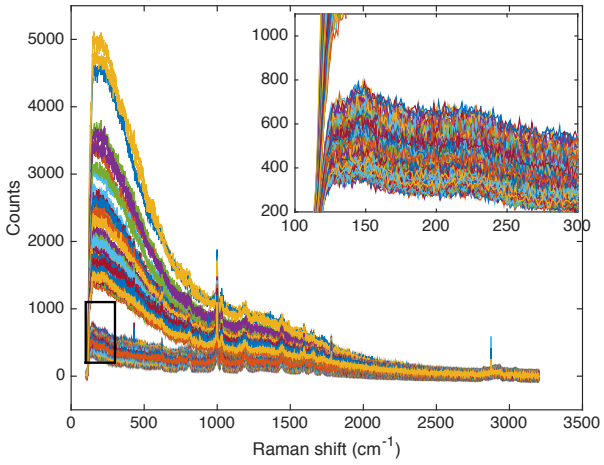

(b)

Figure 20: (a) Raman spectra GNRs embedded ABS filament. The GNRs amplifies Raman scattering of ABS. Inset figure shows the separation between the blank ABS and GNRs embedded ABS Raman spectra. (b) Raman spectra of ABS and DTTCI embedded ABS filaments. Large separation is due to the large quantity of enhancer embedded in ABS filament.

B Detailed Results of Acoustic Classification on Tibial Knee Prosthetic

Tibial Knee Prosthetic Classificatin, Trained with Rectilinear Fill, 60% Density								
	60% Rectilinear Fill		20% Honeycomb Fill			60% Rectilinear Fill		20% Honeycomb Fill	
Index Value	Classification Result	Confidence	Classification Result	Confidence	Index Value	Classification Result	Confidence	Classification Result	Confidence
0	Taz6Tibia_Rectilinear_60_T(0)	132	Taz6Tibia_Rectilinear_60_T(0)	137	43	Taz6Tibia_Rectilinear_60_T(43)	57	Taz6Tibia_Rectilinear_60_T(53)	7
1	Taz6Tibia_Rectilinear_60_T(1)	80	Taz6Tibia_Rectilinear_60_T(1)	19	44	Taz6Tibia_Rectilinear_60_T(44)	70	Taz6Tibia_Rectilinear_60_T(5)	7
2	Taz6Tibia_Rectilinear_60_T(2)	117	Taz6Tibia_Rectilinear_60_T(3)	10	45	Taz6Tibia_Rectilinear_60_T(45)	31	Taz6Tibia_Rectilinear_60_T(55)	8
3	Taz6Tibia_Rectilinear_60_T(3)	108	Taz6Tibia_Rectilinear_60_T(3)	19	46	Taz6Tibia_Rectilinear_60_T(46)	53	Taz6Tibia_Rectilinear_60_T(58)	12
4	Taz6Tibia_Rectilinear_60_T(4)	133	Taz6Tibia_Rectilinear_60_T(4)	18	47	Taz6Tibia_Rectilinear_60_T(47)	28	Taz6Tibia_Rectilinear_60_T(36)	9
5	Taz6Tibia_Rectilinear_60_T(5)	178	Taz6Tibia_Rectilinear_60_T(5)	45	48	Taz6Tibia_Rectilinear_60_T(48)	29	Taz6Tibia_Rectilinear_60_T(17)	10
6	Taz6Tibia_Rectilinear_60_T(6)	61	Taz6Tibia_Rectilinear_60_T(33)	13	49	Taz6Tibia_Rectilinear_60_T(49)	23	Taz6Tibia_Rectilinear_60_T(61)	15
7	Taz6Tibia_Rectilinear_60_T(7)	107	Taz6Tibia_Rectilinear_60_T(12)	9	50	Taz6Tibia_Rectilinear_60_T(50)	41	Taz6Tibia_Rectilinear_60_T(62)	27
8	Taz6Tibia_Rectilinear_60_T(8)	114	Taz6Tibia_Rectilinear_60_T(10)	28	51	Taz6Tibia_Rectilinear_60_T(51)	67	Taz6Tibia_Rectilinear_60_T(63)	15
9	Taz6Tibia_Rectilinear_60_T(9)	189	Taz6Tibia_Rectilinear_60_T(13)	14	52	Taz6Tibia_Rectilinear_60_T(52)	31	Taz6Tibia_Rectilinear_60_T(63)	14
10	Taz6Tibia_Rectilinear_60_T(10)	136	Taz6Tibia_Rectilinear_60_T(13)	45	53	Taz6Tibia_Rectilinear_60_T(53)	23	Taz6Tibia_Rectilinear_60_T(64)	16
11	Taz6Tibia_Rectilinear_60_T(11)	189	Taz6Tibia_Rectilinear_60_T(19)	10	54	Taz6Tibia_Rectilinear_60_T(54)	25	Taz6Tibia_Rectilinear_60_T(66)	33
12	Taz6Tibia_Rectilinear_60_T(12)	194	Taz6Tibia_Rectilinear_60_T(19)	11	55	Taz6Tibia_Rectilinear_60_T(55)	49	Taz6Tibia_Rectilinear_60_T(0)	10
13	Taz6Tibia_Rectilinear_60_T(13)	178	Taz6Tibia_Rectilinear_60_T(16)	72	56	Taz6Tibia_Rectilinear_60_T(56)	31	Taz6Tibia_Rectilinear_60_T(68)	7
14	Taz6Tibia_Rectilinear_60_T(14)	128	Taz6Tibia_Rectilinear_60_T(16)	15	57	Taz6Tibia_Rectilinear_60_T(57)	35	Taz6Tibia_Rectilinear_60_T(68)	17
15	Taz6Tibia_Rectilinear_60_T(15)	204	Taz6Tibia_Rectilinear_60_T(18)	47	58	Taz6Tibia_Rectilinear_60_T(58)	43	Taz6Tibia_Rectilinear_60_T(10)	15
16	Taz6Tibia_Rectilinear_60_T(16)	203	Taz6Tibia_Rectilinear_60_T(15)	14	59	Taz6Tibia_Rectilinear_60_T(59)	49	Taz6Tibia_Rectilinear_60_T(71)	10
17	Taz6Tibia_Rectilinear_60_T(17)	120	Taz6Tibia_Rectilinear_60_T(20)	67	60	Taz6Tibia_Rectilinear_60_T(60)	36	Taz6Tibia_Rectilinear_60_T(71)	83
18	Taz6Tibia_Rectilinear_60_T(18)	147	Taz6Tibia_Rectilinear_60_T(24)	9	61	Taz6Tibia_Rectilinear_60_T(61)	32	Taz6Tibia_Rectilinear_60_T(72)	68
19	Taz6Tibia_Rectilinear_60_T(19)	71	Taz6Tibia_Rectilinear_60_T(27)	10	62	Taz6Tibia_Rectilinear_60_T(62)	31	Taz6Tibia_Rectilinear_60_T(73)	38
20	Taz6Tibia_Rectilinear_60_T(20)	67	Taz6Tibia_Rectilinear_60_T(23)	37	63	Taz6Tibia_Rectilinear_60_T(63)	36	Taz6Tibia_Rectilinear_60_T(74)	14
21	Taz6Tibia_Rectilinear_60_T(21)	99	Taz6Tibia_Rectilinear_60_T(24)	27	64	Taz6Tibia_Rectilinear_60_T(64)	42	Taz6Tibia_Rectilinear_60_T(32)	9
22	Taz6Tibia_Rectilinear_60_T(22)	99	Taz6Tibia_Rectilinear_60_T(32)	12	65	Taz6Tibia_Rectilinear_60_T(65)	46	Taz6Tibia_Rectilinear_60_T(84)	10
23	Taz6Tibia_Rectilinear_60_T(23)	115	Taz6Tibia_Rectilinear_60_T(27)	23	66	Taz6Tibia_Rectilinear_60_T(66)	31	Taz6Tibia_Rectilinear_60_T(84)	10
24	Taz6Tibia_Rectilinear_60_T(24)	70	Taz6Tibia_Rectilinear_60_T(27)	20	67	Taz6Tibia_Rectilinear_60_T(67)	19	Taz6Tibia_Rectilinear_60_T(80)	13
25	Taz6Tibia_Rectilinear_60_T(25)	100	Taz6Tibia_Rectilinear_60_T(25)	11	68	Taz6Tibia_Rectilinear_60_T(68)	18	Taz6Tibia_Rectilinear_60_T(84)	9
26	Taz6Tibia_Rectilinear_60_T(26)	58	Taz6Tibia_Rectilinear_60_T(30)	20	69	Taz6Tibia_Rectilinear_60_T(69)	21	Taz6Tibia_Rectilinear_60_T(30)	7
27	Taz6Tibia_Rectilinear_60_T(27)	41	Taz6Tibia_Rectilinear_60_T(32)	19	70	Taz6Tibia_Rectilinear_60_T(70)	34	Taz6Tibia_Rectilinear_60_T(82)	8
28	Taz6Tibia_Rectilinear_60_T(28)	49	Taz6Tibia_Rectilinear_60_T(33)	14	71	Taz6Tibia_Rectilinear_60_T(71)	70	Taz6Tibia_Rectilinear_60_T(5)	16
29	Taz6Tibia_Rectilinear_60_T(29)	60	Taz6Tibia_Rectilinear_60_T(34)	44	72	Taz6Tibia_Rectilinear_60_T(72)	96	Taz6Tibia_Rectilinear_60_T(10)	11
30	Taz6Tibia_Rectilinear_60_T(30)	93	Taz6Tibia_Rectilinear_60_T(35)	11	73	Taz6Tibia_Rectilinear_60_T(73)	46		
31	Taz6Tibia_Rectilinear_60_T(31)	78	Taz6Tibia_Rectilinear_60_T(35)	34	74	Taz6Tibia_Rectilinear_60_T(74)	36		
32	Taz6Tibia_Rectilinear_60_T(32)	60	Taz6Tibia_Rectilinear_60_T(10)	10	75	Taz6Tibia_Rectilinear_60_T(75)	38		
33	Taz6Tibia_Rectilinear_60_T(33)	53	Taz6Tibia_Rectilinear_60_T(38)	12					