

XCheck: Verifying Integrity of 3D Printed Patient-Specific Devices via Computing Tomography

Zhiyuan Yu, Yuanhaur Chang, Shixuan Zhai, Nicholas Deily, and Tao Ju,
Washington University in St. Louis; XiaoFeng Wang, *Indiana University Bloomington*;
Uday Jammalamadaka, *Rice University*; Ning Zhang, *Washington University in St. Louis*

<https://www.usenix.org/conference/usenixsecurity23/presentation/yu-zhiyuan-xcheck>

This paper is included in the Proceedings of the
32nd USENIX Security Symposium.

August 9-11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.

XCheck: Verifying Integrity of 3D Printed Patient-Specific Devices via Computing Tomography

Zhiyuan Yu[†], Yuanhaur Chang[†], Shixuan Zhai[†], Nicholas Deily[†],
Tao Ju[†], XiaoFeng Wang[§], Uday Jammalamadaka[‡], Ning Zhang[†]

[†] Washington University in St. Louis

[§] Indiana University Bloomington

[‡] Rice University

Abstract

3D printing is bringing revolutionary changes to the field of medicine, with applications ranging from hearing aids to regrowing organs. As our society increasingly relies on this technology to save lives, the security of these systems is a growing concern. However, existing defense approaches that leverage side channels may require domain knowledge from computer security to fully understand the impact of the attack.

To bridge the gap, we propose XCheck, which leverages medical imaging to verify the integrity of the printed patient-specific device (PSD). XCheck follows a defense-in-depth approach and directly compares the computed tomography (CT) scan of the printed device to its original design. XCheck utilizes a voxel-based approach to build multiple layers of defense involving both 3D geometric verification and multivariate material analysis. To further enhance usability, XCheck also provides an adjustable visualization scheme that allows practitioners' inspection of the printed object with varying tolerance thresholds to meet the needs of different applications. We evaluated the system with 47 PSDs representing different medical applications to validate the efficacy.

1 Introduction

3D Printing as Life-saving Technology: Additive manufacturing (AM), or 3D printing, has brought revolutionary changes to a wide range of areas in medicine, from medical instruments to regenerative tissue engineering, from day-to-day clinical practices to biomedical and pharmaceutical research. The flexibility of 3D printing enables the low-cost manufacturing of instruments and devices with complex geometries that are matched to a patient's anatomy. Recent advances in medical 3D printing allow for the fabrication of artificial kidneys, which have been used in real life to save children [62]. While applications such as artificial organs and organ-on-chip are still in their infancy, other medical applications such as prosthetics and hearing aids are widely deployed. Over 10,000,000 people are wearing 3D printed hearing aids and 97% of all

hearing aids globally are being created using AM [48]. The global 3D printing healthcare market size was valued at \$973 million in 2018 and is projected to reach \$3.7 billion by 2026, growing at a compounded annual growth rate of 18.2% [49].

Security Concerns from FDA and Real-world Threat: With the increasing reliance on computing in medical treatment, there are growing concerns about the trustworthiness of modern networked medical systems, with recent demonstrations of vulnerabilities in knee implant [6], pacemakers [25], and medical CT [43]. Furthermore, the importance of 3D printed medical device security is also highlighted in recent FDA documents [2] and device manufacturers surveys [53]. To further understand the feasibility of the threats, we studied the existing medical 3D printing pipelines and disclosed two real-world vulnerabilities responsibly, with CVEs assigned.

Existing Defenses and Application in Medical Treatment: In recognition of the threat, existing approaches turn to physics to monitor the manufacturing process via side channels such as acoustic [6, 7, 13], optical [5, 23, 36, 55] and magnetic [23]. However, the application domain of medical 3D printing presents unique usability challenges. For instance, the physicians presented with a piece of audio recordings and detection results (malicious or benign) may find it difficult to tell the reason for the PSD being rejected.

Our Approach: To address this limitation and further empower the verification schemes with improved usability, we propose to check and visualize the printed product to provide an additional layer of protection. Leveraging a unique capability widely available to medical establishments - CT imaging, we design and implement a CT-model crosschecking system named XCheck. While the use of CT for quality control has been studied before [16, 19, 58], existing tools either require manual efforts from 3D printing experts or focus only on predictable patterns of naturally introduced defects. Therefore, XCheck is designed to automatically compare the CT scans of the printed PSDs to their original design to detect geometric deviations. XCheck further raises the bar for adversaries aiming to compromise the printing material by checking the

CT density distribution on the PSD. In clinical settings, CT is a cost-effective method, since the technical fee of a CT scan is often less than \$200 compared to the million-dollar expense for the entire organ transplantation.

Technical Challenges in Security Application of Computational Geometry: The core technical problem behind the verification is 3D shape comparison, which is well studied in the field of computational geometry [12, 64]. However, similar to the application of artificial intelligence in security, computational geometry techniques also require additional security considerations. Existing 3D shape comparison techniques [10, 51, 56] are designed to tolerate tiny discrepancies for robustness. Nonetheless, security verification often has to reveal all deviations. XCheck takes a defense-in-depth approach, and builds multi-layered analysis from voxel matching (volumetric pixels in three-dimensional CT scans) to X-ray absorption. Based on the observation that attacks have to alter existing voxels, we propose the voxel-based analysis to detect adversarial manipulations that deviate more than twice the resolution of CT, and further prove the security guarantee of the detection. The second layer is a material validation mechanism that takes advantage of the unique X-ray properties of different materials. To overcome the uncertainty in voxel values, the distribution features are used instead. While this validation may not be bulletproof, it does significantly raise the bar for the adversary, requiring expertise in material science to evade detection. The third layer aims to improve defense against adaptive stealthy attacks meant to hide under the CT resolution. We develop a ray-based volumetric detection that analyzes the subspace volume to capture the stealthy manipulations with impacts on the three-dimensional characteristics of the device, but are too small to be identified by voxel analysis. Lastly, inspired by the gamma analysis from radiotherapy, we propose a new clinical risk quantification method to determine the risk associated with a printed device. To enhance the explainability of the verification process, XCheck also provides interactive visualization of the deviations to allow physicians to make better-informed decisions.

Evaluation and Findings: Our evaluation is designed to cover a wide spectrum of medical applications and a diverse set of attacks. Instead of evaluating a few models in a case study manner, we printed 47 PSDs (8 benign models, 39 malicious cases) from 4 distinct medical printing domains (including lung-on-a-chip in pharmaceutical research, orthopedic screw in implant, dental guide in surgical instruments, and bone scaffold in tissue engineering). We also worked with medical printing experts to develop realistic yet stealthy attacks that can lead to patient harm. Altogether, we developed 39 malicious cases at varying scales on both geometric and material properties of the baseline models for representative medical applications. The consequences of these attacks vary from minor discomfort to life-threatening conditions for patients. The experiment shows that our technique is able to

detect and visualize 37 out of the 39 malicious cases.

Contributions: To the best of our knowledge, XCheck makes the first step towards automated support for medical practitioners to verify the physical integrity of 3D printed medical devices. Our contributions are outlined as follows:

- Building on top of our survey on emerging threats for medical applications of 3D printing, we propose XCheck, a highly effective and mostly automated integrity checking system. Unlike the existing works that monitor the printing procedure, XCheck utilizes medical imaging to verify the integrity of the product, by comparing the CT images of the printed PSDs to their original designs at mm/sub-mm scale.
- To meet the security requirements for physical integrity verification, we propose new security-oriented CT analysis techniques to enable both visualization and quantification of adversarial manipulations on geometry and material. Theoretical proofs are developed to show the detection bounds of the proposed system against adaptive attackers who has knowledge of the system, which has never been done before.
- We evaluated XCheck on 47 printed PSDs, including different categories of adversarial manipulations on four representative types of PSD from different fields of medical applications. Our experiment showcased the effectiveness of our approach, and we have released the source code, PSD models, and CT scans¹.

2 Medical Application of 3D Printing

The utilization of 3D printing in medical applications has been fueled by the unique capability of additive manufacturing to create customized devices using a patient’s medical images. Similar to their biological counterparts, many PSDs involve complex porous structures and tortuous internal channels that would be challenging to produce with traditional methods.

Diverse Medical Printing Techniques and Applications: 3D printing has revolutionized many areas of medicine, from orthopedic and otolaryngology to cardiac vascular and oncology. Due to its wide applicability to medical applications, there is also a large variety of manufacturing techniques to fit the mechanical and physiological requirements of individual patients. The most common printing techniques for medical applications include fused deposition modeling (FDM), powder bed fusion (PBF), multi jet fusion (MJF), stereolithography (SLA), and liquid-based extrusion.

3D printing is widely deployed for medical applications, as shown in Figure 2. 3D printed medical instruments are relatively mature, and are deployed from diagnosis to treatment. For diagnosis, patient-specific anatomical models are

¹<https://3dxcheck.github.io/>.




Figure 1: General workflow of medical printing.




Figure 2: Medical applications of additive manufacturing.

commonly printed via FDM or SLA using a patient's medical images. By providing better visualizations of the patient-specific anatomy, these printed models not only enable more accurate diagnoses [21], but also offer surgeons a unique opportunity for in-depth preoperative planning beyond organ measurements. Geometric accuracy is a critical aspect of these models. For treatments, 3D printed surgical guides are used in various procedures, such as total knee replacement surgeries or dental operations. They are commonly used to assist in drilling angles [15] in the operations, therefore the exact geometry must precisely match the anatomy of the individual patients. Due to their relatively low cost and high precision, dental guides are often manufactured using SLA printers with bio-compatible resin. Another area of rapid development is the pharmaceutical industry. 3D printing enables the production of tablets with more than one active substance and allows for different dissolution profiles for individual patients [29].

Tissue engineering is an area that has garnered significant interest [34]. Applications in this area focus on regenerating tissues for individual patients, such as organs or bones. Despite the recent adaption of tissue engineering in real-life clinical cases, the technology is still in its infancy. Organs-on-chip is a recent breakthrough in pharmaceutical research. These chips are 3D cell-cultured microfluidic devices that provide inhabitation for living cells, aiming to reproduce the physiology of tissues and organs on a chip [8], enabling research on organ physiology and disease etiology. Another rapidly emerging area is artificial organs [39]. Many tissues in human organs are unable to regenerate. Transplantation from donors poses risks of rejection and lifelong immunosuppressant-based therapies are often required [42]. Recent research work has demon-

strated the feasibility of regenerating organs, such as bone and cartilage [28], liver [9], kidney [60] and even heart [41]. In these efforts, 3D printing is used to recreate the natural complex cellular support structures, known as extracellular matrices (ECMs). These structures support cell growth in 3D spaces, and their porous morphologies play an important role in cell attachment and proliferation. As a result, geometric accuracy is essential for the performance of the device.

Medical Device Printing Workflow: Medical printing typically includes three main steps, tissue scanning, design customization, and device manufacturing (shown in Figure 1).

1) *Tissue Scanning* - to design the customized device, the structural and pathological characteristics of the patient's body have to be first obtained via medical imaging techniques such as CT and magnetic resonance imaging (MRI).

2) *Design Customization* - in this step, region of interest (ROI) is reconstructed from the patient scan using patient imaging tools such as OsiriX [50], 3D Slicer [47], and Mimics [40]. Design software such as Autodesk Within Medical [4] or Medical Design Studio [1] is then used to design the PSD. Based on individual treatment, the material, such as titanium alloy or soluble polymer, must also be chosen.

3) *Device Manufacturing* - while the diagnosis and device design are often carried out within the hospital perimeter, the fabrication of PSD is often outsourced to dedicated facilities [30] either within the hospital network or with a specialized manufacturer. The printed model also has to go through post-processing, which is highly application-dependent. For implants, the device is typically washed with alcohol and sterilized; for organ-on-a-chip and tissue engineering, the printed scaffolds are cultured with living cells to reproduce the biological functionality of organs, bones, or other tissues.

3 Threat Model and Assumptions

3.1 Adversary Power

Attacker Motivation: The goal of the attacker is similar to the existing works [6, 13, 23, 66] where he/she aims to tamper the 3D printing process that leads to a malformed product. Since PSDs are generally sensitive to human health, a deformed product can bring harm to a patient's well-being.

Adversary Power: In this work, we make the same assumption as previous works [13, 14, 23]. By compromising the printing pipeline remotely, such as slicing software or printer firmware, an attacker may make changes to the received model or G-code, which describes the design and encodes the detailed steps to print an object respectively. Besides, we assume the attacker knows the design of XCheck, which is generally stronger than both practical scenarios and existing work. As such, the attacker may conduct adaptive attacks trying to evade detection, which we analyzed via both experiments (Section 6) and security analysis (Section 7).

Feasibility of Attack Automation: To show the feasibility of scalable remote attacks, we have analyzed and discovered two vulnerabilities in 3D printing tool chains, which were responsibly disclosed under CVE-2021-44961 and CVE-2021-44962. To further show the possibility of generic geometric attacks such as adding internal hollow regions to the devices, we developed a shellcode with the assumption of pre-existing remote code execution vulnerability, to misuse the temperature setting function. This method can be deployed to compromise different types of designs, as demonstrated on a cube and a dental guide shown in Figure 3.

Generalizing the Attack: Using this capability, attackers can make malicious modifications to the printed objects, such as inserting a hollow sphere in the middle, adopting different fill patterns, or swapping materials in a multi-jet printer. As 3D printing technologies continue to evolve, the number of vectors to implement malicious modifications will also increase. However, a key observation is that almost all the attacks we are aware of, regardless of how they are implemented, cause differences in either geometric properties (shape, volume, etc.) or material properties (elasticity, bio-compatibility, etc.).

(1) *Geometric Attacks:* Attacks that cause geometric deviation can result in either exterior surface modification or internal modification or both. Stealthy modifications on the internal of the PSD can often be more challenging to detect than the external modifications.

(2) *Material Swapping Attacks:* Many advanced 3D printers house different materials using different storage tanks or filaments, a remote adversary can swap materials within the printer. In these attacks, an attacker swaps the material to alter the mechanical or physiological characteristics of the PSD.

3.2 Defender Assumption and System Goal

System Goals: Instead of blocking every possible attack vector of different manufacturing techniques, XCheck focuses on the verification of the final product with the goal of detecting geometric attacks and material swapping attacks.

Assumptions on Defender: Different from the previous work that assumes a trustworthy monitoring infrastructure at the manufacturing site, we make the assumption that the defender has access to CT imaging. Besides, we follow ex-




Figure 3: Adversarially modified cube and dental guide models. The manipulations are stealthy since they maintain overall outlook with significantly different internal geometry.

isting works [5–7, 13, 23, 55] and assume the original design of the PSD is available to the defender, which will serve as the ground truth used by XCheck. In practice, such golden designs can be authorized by the physician, who will also have to be part of the decision in the acceptance of the manufactured device. We also assume that the PSD can be scanned by CT imaging since medical imaging is frequently used for diagnostics. The operating environment of XCheck is, therefore, the same as the medical imaging environment.

4 Related Work

The existing defenses are procedure-based methods [5–7, 13, 23, 55] focusing on monitoring the manufacturing process. XCheck is a product-based method that verifies the finished product, as shown in Table 1. The discussions in the rest of the section will focus on attack detection and quality assurance.

Existing Attack Detection Methods. The existing attack detection methods can be categorized as vision-based and side-channel-based. Vision-based verification compares recorded images of the production process to detect deviations [5, 55], but are limited in accuracy and constrained by the ambient light. The use of side-channel signals for attack detection has drawn significant interest due to its potential to analyze the physical process. The side-channel signals allow the reconstruction of physical events during the printing process for verification [6, 7, 13, 23]. Chhetri *et al.* [13] were the first to adopt analog emissions to detect compromised objects during the printing process. However, leveraging a single channel lacks reliability in verification, and the printing material is not validated. To address these two limitations, Bayens *et al.* [6] proposed a three-layered defense-in-depth verification, where acoustic and vibrational signals are collected to reconstruct and verify geometric shapes, and a nanoparticle embedding procedure is used to validate the material. However, the above defenses face new challenges for the verification of medical devices, including channel capacity, sensitivity to environmental noise, technology-specific physical channels, requirement for pre-recorded data, as well as limited material validation.

Existing Quality Assurance Methods. In medical 3D print-

Table 1: Comparison among existing work on defenses of 3D printing

Category	Technique	Vector	Target		Fidelity		Env. Limitations			Reference
			FDM	Others	Att.	Mat.	Noi.	Mod.	Pre.	
Procedure-based	Vision-based	Digital photos	✓		cm*		✓	✓	✓	[5]
		Digital photos	✓		N/A		✓	✓		[55]
	Side-channel-based	Acoustic Signature	✓		cm		✓	✓	✓	[7]
		Printing audio	✓		mm		✓	✓	✓	[13]
Product-based	Model-based	Audio & Vibration	✓		mm ³ *	✓	✓	✓	✓	[6]
		Video & Vibration & Magnet	✓		mm		✓	✓	✓	[23]
Product-based	Model-based	X-ray reconstruction	✓	✓	sub-mm	✓				XCheck

Other printing techniques include SLA, PBF, MJF, PolyJet. *Indicates the precision level is estimated from the referred paper

Att.=Attack Scale, Env.=Environment, Mat.=Material Validation, Noi.=Noise Sensitive, Mod.=Modification on Environment, Pre.=Pre-recorded Data

ing, quality assurance is essential for safety. Existing work in quality assurance generally checks specific important features such as edge shapes, surface roughness, and protrusions rigidity [3, 11, 20, 31, 54, 57, 58]. They all assume the natural occurrence of manufacturing defects, and focus on the examination of features that can uncover such defects. However, adversarial manipulations can manifest as arbitrary modifications to the design, rendering these tools ineffective. For instance, we followed existing work [20, 61] and measured pore size distribution as a metric for porosity quality assurance. In the malicious bone scaffold model where a region of the porous structure of the bone scaffold was filled solid, the distribution did not change significantly, since the removal of a portion of the pores following the same distribution did not affect the size distribution of the remaining pores.

5 XCheck Design

CT scanners image objects with X-ray beam arrays to observe the energy absorption of the object at various angles. The 2D tomographic images are then merged to produce 3D volumetric images comprising voxels. XCheck verifies the integrity by comparing the CT scan with the original design.

5.1 XCheck Overview

XCheck takes a defense-in-depth approach towards physical integrity verification, from geometry verification via volumetric model comparison to material verification via CT number analysis. As shown in Figure 4, XCheck consists of detection mechanisms that examine both the geometry and material of the 3D printed device and a gamma-analysis-based risk quantification method. The proposed system starts with CT model registration that aligns the CT model with the design model, followed by three layers of defense. The first layer of defense is voxel-based analysis, which verifies that all voxels of the device match the design instead of just descriptive features. To understand its security property, we also develop a formal proof of the bounds. The second layer leverages unique X-ray properties of different matter to verify the printing materials via CT number analysis. A ray-based method is proposed as a fail-safe to defend against the adaptive attacker who attempts to inject faults within the bounds of voxel-based analysis. It




Figure 4: High-level workflow of XCheck.

leverages the invariant volume captured by ray-based subspace to detect stealthy manipulations. The gamma analysis combines the analysis results from the previous layers into a clinical risk quantification score.

5.2 CT - Model Registration

The digitally recovered printed device will need to be aligned with the original design model for shape comparison. However, directly adapting existing algorithms for alignment has high performance costs, due to the complexity of many medical devices. For example, a bone scaffold contains more than 10K points in the CT-scan point cloud. To solve this, we adopt Lloyd’s algorithm [32] to approximate the original model with evenly sized polygons, for use in the alignment stage.

To avoid deforming the geometric features during the alignment stage, we make use of rigid plus scaling registration to coherently change the size and placement of the source model so that it aligns with the target model. This process aims to find a transformation \mathcal{T} to linearly pull the k th point on the source model y_k towards the corresponding point x_k on the target model, and can be expressed as:

$$\mathcal{T}(y_k, \mathbf{R}, s, \mathbf{t}) = s\mathbf{R}y_k + \mathbf{t}, \text{ s.t. } \mathbf{R}^T \mathbf{R} = \mathbf{I}, \det(\mathbf{R}) = 1 \quad (1)$$

where $s \in \mathbb{R}^+$ represents uniform scaling, $\mathbf{R} \in SO(3)$ is the rotation matrix, and $\mathbf{t} \in \mathbb{R}^3$ is the translation vector. The goal of registration is to find a transformation $\theta = (s, \mathbf{R}, \mathbf{t})$ that minimizes the misalignment.

Two-stage registration: Large degrees of degradations such as occlusion and rotation can often cause registration to fail. To solve this problem, a two-stage registration process is introduced. The first stage is based on *Fast Global Registration*

(FGR), which aligns two point sets by finding the correspondence with calculated Fast Point Feature Histogram (FPFH) features [68]. *Coherent point drift* (CPD) [44] is then used to refine the alignment, due to its ability to handle noisy point clouds from the reconstructed model.

Anatomical landmark registration: Anatomical landmarks are a set of biologically meaningful points in an organism, which play an important role in medical diagnosis and treatment. As such, the anatomical accuracy of landmark features is an important metric for many PSDs [27]. XCheck also offers a point set registration for optimizing alignments based on landmarks, with which the clinician can specify the regions of important landmarks to prioritize the alignment upon.

5.3 Voxel-based Geometry Attack Detection

Challenges in Using Existing 3D Model Comparison Techniques for Integrity Verification: Existing shape comparison techniques aim to distinguish 3D objects using a set of well-tuned features. For example, 3D comparison techniques in [12, 64] use projection-based representation vectors as features to tolerate minor discrepancies, which is important for the application of shape retrieval. However, these small deviations can have safety implications for PSDs and are important for security verification. As a result, the direct application of existing 3D model comparison methods is insufficient due to this fundamental difference.

Within this context, detecting geometry manipulations requires new designs of comparison algorithms, which face two major challenges. First, any type of imaging system including CT, is an imperfect translation from the physical device to the digital reconstructed model. As a result, there is no one-to-one mapping of geometry elements (e.g., vertices, surfaces) between the design model and the CT model. Second, meshes comprising surfaces are the most common format in describing models in 3D printing, but they are insufficient in security verification since mesh structure disregards the volumetric features. We propose to approximate both models via voxelization and analyze their displacements, which we name as *voxel analysis*. Building on the intuition that geometry attacks require modifications on the voxel constitutions, we develop added voxel analysis and missing voxel analysis to detect all deviations greater than a bounded constant. To understand the theoretical property of the protection, we also developed a proof of voxel analysis detailed in Section 7.

Our Approach - Differential Voxel Analysis: To address the challenge stemming from the lack of one-to-one mapping of geometry elements, we propose to use the Euclidean distances between each voxel centroid in the base model and the corresponding closest voxel centroid on the compared model as a metric. Depending on if the reconstructed model or the original design model is used as the base, the measured deviation can uniquely identify different types of manipula-

tions. When the reconstructed model is used as the base, the calculated distance is effective at identifying added voxels, thus *added voxel analysis*. However, when the original model is used as the base, the calculated distance is effective at revealing missing voxels, thus *missing voxel analysis*. In added voxel analysis, voxels in the reconstructed model are used to discover discrepancies. For the i^{th} point on the reconstructed model, its distance value ρ will be calculated as the distance between the voxel centroid p_i^r and the least-distanted voxel in P_o , where P_o is the collection of voxels of the original model, given as: $\rho(p_i^r, P_o) = \min\{\rho(p_i^r, y) : y \in P_o\}$, where detection is characterized by a ρ greater than the threshold determined by individual clinical treatment. When the voxels are taken from the reconstructed model, voxels in added regions in the object will fail to match the voxels in the original design model, leading to the discovery of adversarial manipulations. For the missing voxel analysis, the original model is used as the base model, and the reconstructed model is used as the compared model. The formulation is similar to added voxel analysis. Since voxels in the original model are used as the base, it is effective in discovering missing voxels.

Security Guarantee: Intuitively, malicious modifications have to create and/or delete volumetric regions to alter the geometry, which in turn creates new voxels that will trigger the alarm in the analysis system. Since XCheck does not make assumptions on the two shapes and compares all the voxels in the model, it detects all voxel deviations, similar to a byte-by-byte file comparison. However, since the granularity of the reconstructed model is directly related to the fidelity of the CT-scan, XCheck is also limited by the scanning capabilities, more precisely 2ϵ , where ϵ is the CT scan resolution. More in-depth discussions can be found in Section 7.

5.4 Material Validation

Materials used in medical printing not only impact the mechanical properties of the finished PSD, but also the chemical and physicochemical properties. As such, it can be an effective attack vector. Existing work [6] has explored the possibility of adding nanoparticles to filaments for identification. We took a different approach that leverages the principle of the attenuation of electromagnetic rays by different materials to identify unwanted materials in a print.

This principle is used in medical imaging to distinguish between different tissues such as bones and muscles. Each voxel has an intensity value, quantified using the *Hounsfield scale*, which is referred to as *CT number*. This number represents the X-ray attenuation coefficient, based on which the Hounsfield Unit (HU) scale is a linear transformation with water and air defined at 0 and -1000 [17]. Ideally, different materials should exhibit different HU values, but it has been found in medical diagnoses [35] that HU values of the same kind of tissue can have a wide range of readings due to uneven physical densities at different locations. Therefore, X-ray absorption is impacted

by physical density (geometry) and atomic mass (material). While such variation exists in medical image visualizations, the focus is often on enhancing the differences among different types of tissues with prior knowledge of human anatomy, such as muscle (HU: 35 to 55) vs bone (Cortical HU: 1800 to 1900) [46]. On the other hand, for material verification, our goal is to minimize the possibility of undetected adversarial insertions of alternative materials into the PSD. To measure the differences between voxel distributions across the entire printed device, we use *kernel density estimation* (KDE).

The distribution $\hat{f}(x)$ of a set of HU values $\{x_1, x_2, \dots, x_k\}$ from the CT scan can be estimated as:

$$\hat{f}(x) = \frac{1}{kb} \sum_{i=1}^k K\left(\frac{x - x_i}{b}\right), \quad (2)$$

where K is the Gaussian kernel function and b is the smoothing parameter. Since the same type of PSDs generally shares similar geometries, differences in materials usually lead to different distributions of HU values. For quantification, we further extract three features: expectation, highest density, and its corresponding HU value to represent the material.

Experimental Verification: We experimentally verify that our proposed technique can distinguish common printing materials using three types of medical devices (organ-on-chip, dental guide, orthopedic screw) with the seven most common materials including a biocompatible resin. More details for the experiments can be found in Section 6.

5.5 Ray-based Adaptive Attack Detection

Need to Defend Against Adaptive Attacker: In voxel analysis, each voxel is represented by its centroid in the calculation of the distance to its matching point. When voxel analysis is ultimately comparing two point clouds, it disregards the geometric context information, such as the surrounding shapes and voxel connectivity. Exploiting this limitation, an adversary may craft stealthy attacks where the deviation is below the CT scanning limit while preserving the voxel features, but these small deviations will aggregate leading to significant differences in the volume of the PSD to harm the patient. The bone scaffold attack we developed for evaluation is an example of stealthy manipulations, where pore radius is deliberately increased, but the magnitude is carefully planned such that individually they are almost always within the bounded regions, however, the aggregated impact cannot be overlooked.

Our Approach - Ray-based Volumetric Attack Detection: To solve this challenge, we designed a *ray-based volumetric analysis* to detect divergence in geometric volume. The concept of the ray-based feature was first proposed by Vranić *et al.* [65] in the context of information retrieval to describe 3D shapes by measuring the extent of an object in given directions, i.e. along defined rays. However, this initial design is limited in three aspects for device integrity verification.

First, current ray-based approaches place emphasis only on the outer envelope of an object and discard the internal structures. To address this, we modify the method to capture the innermost envelopes. However, it is still ineffective in detecting attacks that are located behind the first layer of internal surfaces around the model centroid. To solve this problem, we emanate rays from each voxel instead of the geometric centroid, which is designed to capture geometric features within a small region around each point. Lastly, when rays are cast with random horizontal and vertical angles they do not result in uniformly distributed points in space (such as a spherical surface). Therefore, we use the Fibonacci lattice [18] to create uniformly spaced rays.




Figure 5: Visualization of ray-based attack detection.

By placing the rays that capture the geometric features of both models, we further consider the following two questions: (1) how to visualize the volumetric differences, and (2) how to quantify volumetric differences for attack detection. Given an origin, when an emanated ray intersects a model, its length correlates with the magnitude of the inconsistency. To capture the discrepancy in volume, we record all rays that have different intersection coordinates between the reconstructed model and the original model. We can capture the different regions by recording the first intersected point because it describes the extent of the deviation. To reconstruct the volume, we generate an *alpha shape* from the intersection points. By using Delaunay complexes and their filtrations, the alpha shape provides a quantitative method to accurately describe and compute shapes at multiple levels of detail.

To summarize, the key idea of ray-based analysis is illustrated in Figure 5. Selected points of origin for rays define the collection of subspaces to verify, for each subspace, the measured volume is then checked against the original design. The selection of points can be random, uniformly distributed within the model space, or voxels discrepancies above a specific threshold, providing different levels of protection.

5.6 Gamma Analysis for Aggregated Risk

While voxel-based analysis and ray-based analysis offer complementary mechanisms to measure geometric inconsistencies, both methods rely heavily on the strong expertise of users to understand the impact of attacks. To address this limitation,

we build on the concept of *gamma analysis* from radiotherapy [37] to aggregate measurements from different analyses of the PSD into a single number describing the clinical risk of the medical device being compromised.

Gamma analysis is an approach used by the medical community to aggregate clinical factors into a single metric [37]. To quantify the integrity of a printed PSD, there are four key parameters, each individually capturing a unique attack vector. First, scaling in registration (γ_S) can be used as an attack vector. Since the registration aims to minimize misalignments between two models, the reconstructed model can be linearly scaled to better align with the source model. An adversary can create a geometrically matched device at a different scale than the original model, while the registered reconstructed model and original model would still match. To defend against this attack, the scaling factor is captured as $S = |1 - s|$, where s is the scaling parameter in alignment. Geometric deviations in the voxel analysis (γ_D) are previously discussed, and the key idea is to use the maximum voxel deviation to describe this factor. This is because if there is one instance of deviation above the threshold, it is unlikely the device can be accepted. For volumetric differences in the ray-based analysis (γ_V), it is designed to capture V as the percentage of device volume that shows mismatched geometry. Lastly for material, the deviation M is described by the expectation, peak density, and the corresponding HU value in the estimated HU value distribution. Incorporating the above elements, the aggregated risk score Γ_{PSD} for the PSD can be formulated as:

$$\Gamma_{PSD} = \sqrt{(\gamma_S)^2 + (\gamma_D)^2 + (\gamma_V)^2 + (\gamma_M)^2} \quad (3)$$

Intuitively, this risk score is an aggregation of the four risk factors. For each risk factor p , γ_p is defined as $\gamma_p = \max\{p, C_p\}/C_p$, where p can be S, D, V or M , while C_p is the acceptable threshold criteria for p . For the individual patient and the PSD, C_p depends on clinical factors such as patient age, sex, and the purpose of the PSD. For example, the scaling factor S is critical for an implant to ensure patient fit, but it is not as important in an anatomical model. The key idea behind the γ_p is that if any deviation within its acceptable threshold, its gamma value would be a constant value of 1, however, the more it deviates beyond its acceptable threshold, the more dominant that factor is in the gamma analysis.

Therefore, a standard object should pass the *gamma analysis* with all four factors $\gamma_S, \gamma_D, \gamma_V$ and γ_M equal to 1, which sets the acceptable threshold for the composite value Γ_{PSD} as 2. Calculated values above this threshold reveal the existence of an attack, and a higher gamma value indicates a higher attack probability or a larger magnitude of the manipulation.

5.7 Adjustable Visualization

Proper visualization is a key element that empowers users to understand the verification results in an intuitive way. Using




Figure 6: User interface of XCheck with/without colormap.

the dental surgical guide in Section 5.3 as an example, after registration and downstream analysis, XCheck provides users with an interface as shown in Figure 6. Several customizations can be applied to aid users in the verification process. The *Colormap* checkbox allows users to apply (or remove) a colormap, which normalizes a color gradient and paints individual voxels based on their calculated Euclidean distance with the matching voxel on the compared model. We also design rendering based on color psychology, where we use blue to render voxels with shorter distances (i.e., safer voxels) and red to render voxels with larger discrepancies (i.e., attacks more likely to happen). Users can also control different aspects of the visualization using the following sliders. (1) *Opacity*: filters out voxels below a certain distance threshold; (2) *Colormap Upper Bound*: paints all voxels with distance above the upper bound red, then normalizes and assigns colors with the new bounds; (3) *Colormap Lower Bound*: paints all voxels with distance below the lower bound blue, then normalize and assign colors with the new bounds; and (4) *Point Size*: adjust voxel size ranging from 1 to 10. By customizing these parameters, users can identify whether a region of interest has been manipulated.

6 Experiments and Evaluations

Evaluation and Attack Design Rationale: Our evaluation is designed to analyze the effectiveness of XCheck in: (1) identifying malicious geometric modifications on devices, (2) detecting adaptive adversarial manipulations that attempt to hide under the scanning resolution, (3) distinguishing the existing commercially available printing materials, and (4) using gamma analysis to quantify potential risks of compromised devices. Instead of evaluating on a few models in a case study manner as done in existing work, we conducted a larger-scale experiment with 47 PSDs covering a variety of medical applications and attack types. Out of the printed models, 8 were benign models, 30 underwent geometric modifications while retaining the original material, and 9 underwent material modifications while maintaining the original geometry. One SLA

printer (ELEGOO Mars UV Photocuring 3D Printer) and one medical CT (Siemens Vision) scanner were used in the process of manufacture and verification. Some of the models are shown in Figure 7.



Figure 7: The printed and scanned models.

We selected four medical applications for evaluation as they uniquely represent different domains. The orthopedic screw represents implant, the bone scaffold represents tissue engineering application, the dental guide represents surgical instrument, and the lung-on-a-chip represents pharmaceutical research. The orthopedic screw is used to evaluate malicious modifications on surfaces since subtle modifications to its geometric surfaces can significantly vary the mechanical properties. The bone scaffold is used to evaluate adapted modifications on internal geometry because of its highly complex porous structure. To evaluate material attacks, we performed tests on the orthopedic screw, lung-on-a-chip, and dental guide models, and quantified the relative sensitivity of each model to changes in material properties of mechanical strength, elasticity, and biocompatibility.

Similar to existing works [5, 6, 23], we have to develop our own attacks due to two reasons. First, naive attacks are often easy to detect, and stealthy adaptive attacks can further challenge our defense. Second, attacks with physiological impacts simulate potential threats in the real world and therefore are better validation cases. To make the models and attacks realistic, the patient-customized base models are developed in collaboration with the medical 3D printing lab in the medical school, with some sourcing from the National Institutes of Health (NIH) resource sharing platforms. We used *Blender*, *Cinema3D* with *Proc3durale*, and *PTC Creo* to create the attacked models and sliced them using *Chitobox*, a popular slicer for resin-based printers. We printed our models using an ELEGOO Mars UV Photocuring 3D Printer with a resolution of 0.047mm. The models were further cleaned with 95% ethyl alcohol during post-processing. The scans were performed on a Siemens CT scanner at 0.6mm slice thickness. Resolution along x and y axes is 512×512 pixels of $1 \times 1\text{mm}^2$. Similar to existing works [5, 6, 23], we did not perform experimental comparisons with previous work due to significant differences in targeted attack vector (all vs specific), methodology (CT vs side-channel), precision (arbitrary changes mm/sub-mm level), and printing technology (generic vs FDM).




Figure 8: Minor diameter attack on orthopedic screws. From left to right: 4.52mm, 4.80mm, 5.06mm minor diameter.

Geometry Attack Evaluation: Orthopedic screws are used to evaluate geometric modifications because changes in the key mechanical characteristics, such as minor diameters, cannot be easily captured by existing 3D shape comparison algorithms, which generally discard minor differences. For this model, we implemented four adversarial manipulations: (1) elongating the non-threaded shank, (2) removing a section of thread, (3) enlarging the minor diameter, and (4) enlarging the pitch size. Due to limited space, only (3) is included in this section.

Figure 8 shows the attack that increases the minor diameter of the threads while preserving the major diameter. The base model's minor diameter is 4.52mm, and we change this to 4.58mm, 4.70mm, 4.80mm, and 5.06mm as adversarial modifications. The left three pictures in Figure 8 show the printed screws. As shown, it is difficult to differentiate these models visually. Such modifications can affect the adhesion of these screws and potentially lead to problems such as causing fracturing or bone separation.




Figure 9: Added voxel analysis (AVA) and missing voxel analysis (MVA) on detecting the enlarged minor diameter.

The detection results are listed in Figure 9, where the 4.70mm model is not shown due to its similarity to the 4.80mm model. It can be seen that all thread flank regions are rendered in red, indicating the existence of attacks. With either the reconstructed model or the original model as base, both *added voxel analysis* (AVA) and *missing voxel analysis* (MVA) are able to highlight the modified region. This is because that when the designed model serves as base (MVA), the original volume around the thread pitches are missing on the printed model, while when the CT is used as base (AVA),




Figure 10: The appearance of thumb proximal phalanx bone scaffold during the creation and manufacture processes.

the volume on the enlarged thread pitch appears as an added feature to the original design. This attack triggers both added and missing voxel discrepancies and is detected by both.

Adaptive Geometry Attack Evaluation: XCheck is evaluated against adaptive attacks with the bone scaffold models, on which we specially crafted stealthy modifications that cause minor changes below the scanning resolution of the CT. The pore size within a small region (e.g., radius of 0.25mm) is reduced by 0.14mm , whereas the CT resolution is 1mm . We show that these attacks are effective in evading the voxel-based geometric comparison due to their small magnitude of changes, however, the aggregated volume from such small changes can still harm the patient in unexpected ways.

Bone scaffolds play an important role in *tissue engineering* applications, as they provide support for cell attachment and growth, similar to extracellular matrix (ECM) in native tissues. As with most scaffolds, the pore size, interconnectivity, and shape are key factors. They are important in various tissue development stages, from cellular attachment and motility to nutrient diffusion and waste removal. As a result, small changes in the diameter of the pores or their interconnectivity can lead to severe consequences.

Figure 10 shows the three steps of creating the thumb proximal phalanx bone scaffold. The initial model without porous structure (top left) is generated from a patient model with surface remeshed while preserving the geometry. Then a porous volumetric texture is created using a 3D shading tool called *Proc3durale* to create the scaffold (top right). The printed model is shown at the bottom.

We conducted four adversarial manipulations: (1) adding material in the form of internal solid regions, (2) removing material to leave internal hollow regions, (3) reducing the pore size in certain regions, and (4) rotating the condyle to different extents. Limited by space, we present the results of the third attack due to its stealthiness and potentially high impact on physiology.

Previous research has demonstrated that scaffold porosity and pore size can affect cell attachment and efficacy, mechanical strength, the ability to promote *in vivo* osteoconduction and vascularization, and *in vivo* and *in vitro* cell signaling,




Figure 11: Missing voxel analysis (MVA) and ray-based analysis on detecting the region of reduced pore size.

thus affecting the ability of a scaffold to support new tissue formation [33]. For this attack, four malicious models were designed, each with an internal region that contained modified porous structures. The affected regions in each model had a radius of 0.25mm , 0.75mm , 1.25mm , and 1.75mm respectively. The average pore size in these regions was reduced from $155\mu\text{m}$ to $15\mu\text{m}$, which could significantly impact the nutrient flow in these regions. Such change can have severe consequences in the cell cultures, with decreased cell proliferation and nutrient availability in these regions. Besides pore size, the affected region also plays an important role, since the larger the tampered region, the more likely the cell cultures will be affected during proliferation.

Figure 11 shows that XCheck can detect and visualize modified porous regions as small as 0.25mm . Note that we only show the results for 0.25mm and 1.25mm due to limited space. While users may be able to visualize internal discrepancies in voxel analysis by adjusting opacity, many modifications on complex internal geometries can be difficult to visualize using this method, as shown in Figure 11(a), 11(c). The voxel analysis is only able to locate deviations that exceed 0.75mm , yet it shows only a small portion of the impacted area. It is because there are too many voxels with similar distances, which result in noise and cause the tool to miss the deviations. On the contrary, the ray-based analysis extracts volumetric shapes around the impacted region, which is more informative about how the geometry differs and how much volume is maliciously affected. The results of the ray-based analysis are shown in Figure 11(b), 11(d).

Material Attack Evaluation: We focus on different types of materials available for SLA printers. Similar to previous work [6], we did not test on organic matters, since scaffolding often determines the final organization of tissues, and various challenges in experimenting on live cells. While this may limit us to a smaller set of materials, the collection still gives us a reasonable evaluation of the feasibility of our detection mechanism. Materials commonly used in SLA printers are characterized by three main properties, hardness, elongation, and biocompatibility. We selected three models, orthopedic screw, lung-on-a-chip, and dental guide to test material attacks

that aim to compromise the mechanical strength, elasticity, or biocompatibility of each device respectively. In the following section, the orthopedic screw models were used for case study.

Table 2: Compared materials utilized for SLA printing

Name	Feature	Hardness	Elongation	*Bio
LCD-MG	rigid & tough	80D	5%	✗
LCD-C	casting	80A	24%	✗
LCD-N	nylon-alike	75D	115%	✗
LCD-E	elastic	26A	300%	✗
LCD-G	super clear	70D	5%	✗
SG Guide	biocompatible	70D	7%	✓

*Bio stands for Bio-compatibility.

For the orthopedic screws, mechanical strength is a key attribute. Five different resins, *LCD-E*, *LCD-C*, *SG Guide*, *LCD-N*, and *LCD-MG* were adopted to change the mechanical strength of this model, where their hardnesses are 26A, 80A, 70D, 75D, 80D, respectively.

Using all the orthopedic screws printed in the geometry evaluation as screws with different specs (i.e., length, pitch size, minor diameter, and thread number) but the same material, we aim to measure the deviations of these materials from the standard material. Figure 12 shows the results, where the z-axis is the expectation of distribution, the y-axis is the highest density of the estimated KDE function, and the x-axis is the corresponding HU value of the highest density.

In our experiments, the orthopedic screws varied in length by 1.2cm (22%), pitch size by 0.5cm (18%), number of sections by one (9%), and minor diameters by 0.54mm (12%). Despite the different geometries, they remain closely clustered as shown in Figure 12, as compared to the screws with the same specifications but different materials. We also observed that even though material composition may play a larger role in the HU value and density, geometry also matters. As a result, PSDs with a similar high-level design would exhibit a similar signature in the KDE.

6.1 Gamma Analysis Evaluation

As previously defined, the value of Γ_{PSD} is determined by four factors $\gamma_S, \gamma_D, \gamma_V, \gamma_M$, with each monitoring an attack vector. Therefore, a PSD is considered benign when its Γ_{PSD} value is 2. Since deviations will be recorded if they exceed the acceptable thresholds C_S, C_D, C_V and C_M , these parameters need to be carefully selected based on the individual patient and the medical application. There are two main considerations. (1) The smallest feature deviation across the entire device. For example, the smallest feature on the surface of an orthopedic screw is the threads. In our selected model, the distance between the non-threaded shank and each thread is 0.8mm, the acceptable distance discrepancy C_D can be set to this value. (2) The highest resolution in the manufacturing pipeline. Particularly, the resolution of the CT scanner and




Figure 12: Material clusters - each point shows the expectation, the highest density and the corresponding HU in KDE.

the medical printer impose a physical limitation on how low the acceptable deviation can be set.

Following this rationale, we tested the proposed gamma analysis on each attack implementation. Working with medical experts, we set the acceptable deviation levels according to the device operating environment. The results from the gamma analysis can be found in Table 3 in the Appendix, where we successfully detected 37 out of 39 attacks implemented on the four types of models. Note that our detection results are outcome-oriented, and the system will reveal all the discrepancies greater than the specified value in the policy. PSDs that have not been modified and were properly produced generally exhibit minor deviations due to imprecision in the manufacturing process. In such cases, the deviations are relatively small and therefore won't trigger the detection system. For malicious devices, there have to be volumetric modifications to have meaningful physical world effects and therefore will be detected by XCheck. For modified but not malicious PSDs, the XCheck will also raise an alarm, since there is no mechanism for the machine to determine if the modification will have clinical risks or not. The physician has to make the call in this case.

True Negative, False Negative, and False Positive: The results of XCheck include both false positives (FP) and false negatives (FN). In our evaluation, there were two false negative cases, in which the modification magnitudes were significantly below the CT resolution. Such failures are within our expectations since it is hard to capture attacks smaller than the CT resolution. To gain better accuracy, one can use microCT that offers a higher resolution of $10\mu\text{m}$. A low false positive rate (i.e., high true negative rate) is essential for XCheck, as it is anticipated that the majority of all manufactured devices are benign in practice. In the experiments involving eight benign models, we encountered two possible FPs. The two FP cases occurred in the lung-on-a-chip models where geometry attacks were detected. Upon closer examination, we identified tiny manufacturing defects where the hollow microfluidic channels were not properly molded, likely due to

the resolution limitations of printers. This also shows that our tool may not differentiate between random and malicious modifications, similar to a byte-by-byte comparison-based file integrity verification in the cyber domain.

It is important to note that the low FP rate in our evaluation relies heavily on both hardware capabilities and threshold selection. The printer and scanner used during experiments both exhibit relatively high resolution, which enabled the successful manufacture of devices and accurate digital reconstruction. There were two cases of manufacturing defects in our evaluation due to limitations on printer capabilities, and two cases where XCheck failed to detect the attack due to limitations on imaging capability. Therefore, it is important that appropriate hardware is selected for the corresponding clinical application. Additionally, the selection of the threshold is also a key factor. In the experiment, we worked with medical experts to set the threshold, since the threshold is likely to be different for each clinical application domain. When the thresholds were set to extremely small without consideration for the hardware capability or clinical application domain, the FP rate will increase significantly. While our tool provides automatic analysis and visualization for the verification process, it is ultimately up to the medical professional to make the informed decision on whether to accept the printed PSD or not.

Run-time Efficiency: We also evaluated the run time efficiency of XCheck when verifying each printed PSD. The run time was recorded from the start of registration to the end of gamma analysis, including all the involved analysis techniques. The average run time for the evaluated 47 PSDs is 258.2 seconds. More details are in Table 3 in the Appendix.

7 Security Analysis

Theorem 1 characterizes the security of XCheck.

Definition 1. Let V_M be the set of voxel centroids in voxelized model M , $d(x, y)$ be the Euclidean distance of point x and y , the distance $D(x, M)$ of a point x to a model M is defined as $\min\{d(x, y) | y \in M\}$ (distance to the closest voxel centroid in the matching model). We denote the CT reconstructed model M_c and design model M_d to be ϵ -geometrically bounded, if deviation $\max[Dev(M_c, M_d)] < \epsilon$, where:

$$Dev(M_c, M_d) = \{D(x, M_d) | x \in V_{M_c}\} \cup \{D(y, M_c) | y \in V_{M_d}\}$$

Definition 2. Voxel change magnitude is defined as the minimum displacement between any voxel centroid v_{old} in V_{old} and any point p_{new} in V_{new} , where V_{old} is an existing volumetric model of which a voxel has the shortest distance to any voxels in maliciously crafted new model V_{new} , written as $\min\{d_s(v_{old}, p_{new}) | v_{old} \in V_{old}, p_{new} \in V_{new}\}$.

Theorem 1 (2ϵ -geometrically bounded attacks). *Under the assumption that M_c is an accurately reconstructed model of printed device using computed tomography with resolution*

of ϵ , XCheck guarantees 2ϵ -geometrically bounded matching of the reconstructed model M_c and M_d , therefore detects all geometric attacks with voxel change greater than 2ϵ .

A proof is given in Appendix A.1. Intuitively, Theorem 1 states that XCheck is guaranteed to detect all geometry attacks that cause voxel changes greater than twice the threshold of scanning limit. A key idea of the proof is that the added portions of the geometry will inevitably introduce new voxels to describe the additional volume, and a similar rationale also applies to attacks that remove voxels. As a result, the two models are no longer 2ϵ -geometrically bounded, causing the attack to be detected. Due to uncertainties in the reconstruction, the guarantee can only be made when the range of possible values between the attack voxel and the original voxel do not overlap, thus 2ϵ .

Defending Adaptive Geometric Attacks: While changes bigger than 2ϵ can be revealed, detecting attacks below this threshold in a noisy real-world environment is probabilistic. Knowing the imaging limit of CT and the tolerated deviation of the defender, an attacker can restrict the changes to be slightly below the threshold in an attempt to evade detection using a highly precise (sub-mm) printer. The detection accuracy is further limited by not only uncertainties in imaging but also uncertainties in printing. However, an invariant is that in order to cause effective harm to the patient, the attacker needs to make such attacks over a substantial area to have physiological effects, which in turn increases the chance of being detected. With high-precision images, even small regions have a large number of voxels. For instance, it requires modification on over 200 voxels to create a $1mm^3$ hollow sphere within the bone scaffold model. The stealthy attack also needs to maintain invariant volume among sub-spaces to evade ray-based detection, while remaining undetected by the added and missing voxel analysis. Further analysis with noise distribution model is available in Appendix A.2.

Defending Adaptive Material Attacks: For material attacks, the adversary capability varies significantly between a remote attacker and a physical attacker. Remote attackers can only switch existing materials that are loaded into the printer, therefore the choice is very limited. Most of the time, materials loaded into different containers of the printer present different CT-response profiles, as shown in our experiment on the 7 types of main materials for SLA printing. Therefore they can be detected using XCheck.

With physical access, an attacker can make arbitrary changes to the materials, which makes detection much more challenging due to the diversity of harmful materials. In general, materials often exhibit different X-ray attenuation properties. Even when they have similar HU ranges, their variances are different, such as high impact polystyrene (HIPS) and polylactic acid (PLA) [45]. Additionally, they differ in other properties such as density, hardness, tensile strength, and best printing temperature [59]. Therefore it is difficult to find

materials that exhibit similar CT responses and mechanical properties. Furthermore, many organic toxins are generally fragile under heat or UV light, and are therefore unlikely to survive the manufacturing process. For example, active virus is relatively low resistant to UV light [38]. Moreover, many contaminants such as heavy metals and radioactive elements, will leave significant footprint in CT images. As a result, it is non-trivial to find material that can maintain toxicity through the manufacturing process while keeping similar mechanical and optical properties. Lastly, we would like to emphasize that XCheck likely cannot prevent all types of lethal toxins using only CT scans. However, it significantly raises the bar for attackers who not only need to have access to non-traceable toxic materials and anonymous physical access to the facility to avoid attribution, but also need to possess strong expertise in material science and human physiology.

8 Limitations and Discussions

Cost: Technical fee including the costs directly related to the scan is generally less than \$200 for each CT scan [52]. For regular devices, the scanning cost for each can be even lower, as a batch of multiple devices can be scanned at the same time in practice. On the other hand, treatment costs for PSD implanting, such as organ transplants can cost millions of dollars. Considering the risk to patient lives and relative treatment cost, XCheck is a cost-effective approach.

Dependency on CT Scanning: Even though the methodology in XCheck is independent of the CT scanner used, the performance of the system heavily relies on CT imaging when deployed, since it is the bridge between the actual physical medical device and the digital reconstruction used for verification. For this technology to transition to practice from fundamental research, the following aspects should be further investigated. The first aspect is the accuracy and resolution of the CT scanner. The smallest unit of scanning is captured by the resolution. In the context of security verification, adversarial modifications have to be greater than twice the resolution for XCheck to guarantee detection. Accuracy is less well-defined, for the purpose of discussion, we consider the probability of the CT scan matching the physical objects. In general, the less the uncertainty is, the higher the accuracy is, and the lower the false positive is. The second aspect is the calibration of the CT machine. Calibration is a maintenance process that takes place regularly. The lack of calibration can result in erroneous HU value reports, thereby leading to incorrect patient diagnosis and inaccurate analysis results of XCheck. To mitigate the risks, radiologists and department imaging directors typically establish rigorous protocols that practitioners must follow to ensure the optimal functioning of CT imaging [26]. The third aspect is the scanning orientation of specimens. For patient safety, this has been studied extensively [22, 63] and research indicates high concordance of shape measurement us-

ing medical CT [22], and the shape variation can be controlled under 10 micrometers when the recommended orientation is followed on industrial high-resolution CT machines [63]. In practice, standard clinical examination CT protocols dictate specific object orientations for optimal imaging quality [24]. The last perspective is scannable material. While medical CT imaging is effective in distinguishing tissues, its reliance on the diversity of X-ray attenuation properties of various materials poses some limitations. For example, medical CT machines generally cannot scan metal objects. Besides, special materials may be exploited such that some internal layers of the PSD are not clearly visualized in CT scans. In such cases, the inaccuracy from CT recovery will likely trigger a rejection of the medical device and raise alarm, since the geometry does not match the design. The physician could leverage stronger CT (such as industrial CT) machines to scan metal implants or penetrate the material.

Printers and Scanners Variations: XCheck is designed to inspect the final product, therefore, it is capable of disentangling the production process and the verification process. Even though the separation from the printing process allows XCheck to apply to different printing technology, it also prevents the procedure from taking advantage of the unique physical behaviors of each printer technology. Unlike printer variation, scanner variation has a greater impact on the performance of XCheck. The higher the scanner resolution is, the more accurate XCheck is. The higher the accuracy of the scanner, the better XCheck performs. Fortunately, all medical imaging facilities have periodical maintenance to ensure the optimal performance of the scanner [26].

Scalability of XCheck: XCheck is designed to check each PSD against its design individually, therefore its scalability can be limited especially when the verification involves humans (e.g., physicians) in the loop. After consulting with medical practitioners, we believe the verification process for each PSD can be integrated into treatment/surgery preparation which could take hours, while analysis performed by XCheck typically takes several minutes. In practice, we consider such efforts reasonable in light of the potential risks.

Applications beyond Medical Devices: Outside the realm of PSDs, the core techniques of XCheck can have broad impacts in verifying the integrity of other physical objects that may be of cyber origin, including the medical devices or mechanical components that are not manufactured by 3D printing. As such, it can potentially serve as a general tool for quality assurance that provides a more comprehensive verification. Beyond that, they also have potential applications in other areas. For instance, the techniques could be adapted to augmented reality (AR) for inspection assistance, or be used to compare the shape of the same object but sensed via different sensors (LiDAR and Radar) to prevent injection attacks [67]. Techniques we developed for XCheck is a step forward in enabling its application for security.

Ethical Considerations: We follow the best practice from the community to address ethical considerations. First, we disclosed the vulnerabilities to the vendor and collaborated with them to understand and remediate the issues, thereby reducing the risk of adversaries exploiting them for nefarious purposes. Second, while the high-level idea of attacks is discussed in the paper, implementation is non-trivial requiring expertise in computer security, manufacturing, and human physiology, as well as a significant commitment. Lastly, by making the vulnerabilities known to the community, we aim to shed light on the new threat landscape to motivate effective protection before it manifests in the real world and harms the patients.

9 Conclusion

Based on our analysis of existing medical applications of 3D printing, we proposed XCheck, a cyber-physical cross-checking system that leverages widely available access to medical imaging to verify the integrity of patient-matched medical devices by comparing CT scans with their design models. Building on techniques from computational geometry, we developed new 3D shape comparison techniques along with its security proof to address new challenges due to their application in security. We experimentally evaluated our system on 47 PSDs to explore the feasibility and limitations.

Acknowledgment

We thank the reviewers for their feedback. This work is supported in part by US National Science Foundation under grants CNS-1916926, CNS-2038995, CNS-2154930, and CNS-2238635, and by Army Research Office under contract W911NF-20-1-0141.

References

- [1] Medical design studio - anatome. <https://www.anatome.com/medical-design-studio>, Jan 2017.
- [2] Discussion paper: 3d printing medical devices at the point of care. <https://www.fda.gov/medical-devices/3d-printing-medical-devices/3d-printing-medical-devices-point-care-discussion-paper>, Dec 2021.
- [3] Mohammad S Alsoufi, Abdulrhman E Elsayed, et al. Surface roughness quality and dimensional accuracy—a comprehensive analysis of 100% infill printed parts fabricated by a personal/desktop cost-effective fdm 3d printer. *Materials Sciences and Applications*, 2018.
- [4] Autodesk. Medical 3d printing and orthopedic implant design software. <https://www.autodesk.com/products/within-medical/overview>, Nov 2015.
- [5] Felix Baumann and Dieter Roller. Vision based error detection for 3d printing processes. In *MATEC web of conferences*. EDP Sciences, 2016.
- [6] Christian Bayens, Tuan Le, Luis Garcia, Raheem Beyah, Mehdi Javanmard, and Saman Zonouz. See no evil, hear no evil, feel no evil, print no evil? malicious fill patterns detection in additive manufacturing. In *26th USENIX Security Symposium*, pages 1181–1198. USENIX Association, August 2017.
- [7] Sofia Belikovetsky, Yosef A Solewicz, Mark Yampolskiy, Jinghui Toh, and Yuval Elovici. Digital audio signature for 3d printing integrity. *IEEE Transactions on Information Forensics and Security*, 2018.
- [8] Sangeeta N Bhatia and Donald E Ingber. Microfluidic organs-on-chips. *Nature biotechnology*, 2014.
- [9] Sangeeta N Bhatia, Gregory H Underhill, Kenneth S Zaret, and Ira J Fox. Cell and tissue engineering for liver disease. *Science translational medicine*, 6(245):245sr2–245sr2, 2014.
- [10] Silvia Biasotti, Andrea Cerri, Alex Bronstein, and Michael Bronstein. Recent trends, applications, and perspectives in 3d shape similarity assessment. In *Computer graphics forum*, volume 35, pages 87–119. Wiley Online Library, 2016.
- [11] Irene Buj-Corral, Xavier Sánchez-Casas, and Carmelo J Luis-Pérez. Analysis of am parameters on surface roughness obtained in pla parts printed with fff technology. *Polymers*, 13(14):2384, 2021.
- [12] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. On visual similarity based 3d model retrieval. *Comput. Graph. Forum*, 2003.
- [13] Sujit Rokka Chhetri, Arquimedes Canedo, and Mohammad Abdullah Al Faruque. Kcad: kinetic cyber-attack detection method for cyber-physical additive manufacturing systems. In *2016 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–8. IEEE, 2016.
- [14] Ang Cui, Michael Costello, and Salvatore Stolfo. When firmware modifications attack: A case study of embedded exploitation. 2013.
- [15] A. Dawood, B. Marti, and V. Sauret-Jackson. 3d printing in dentistry. *British Dental Journal*, 219:521–529, December 2015.
- [16] Leonardo De Chiffre, Simone Carmignato, J-P Kruth, Robert Schmitt, and Albert Weckenmann. Industrial applications of computed tomography. *CIRP annals*, 63(2):655–677, 2014.

[17] Tami D DenOtter and Johanna Schubert. Hounsfield unit. In *StatPearls [Internet]*. Treasure Island (FL): StatPearls Publishing, Jan 2022.

[18] R Dixon. Mathographics. basic blackwell limited, 1987.

[19] Anton du Plessis, Stephan Gerhard le Roux, Gerrie Booysen, and Johan Els. Quality control of a laser additive manufactured medical implant by x-ray tomography. *3D Printing and Additive Manufacturing*, 2016.

[20] Anton du Plessis, Stephan Gerhard le Roux, and Anina Guelpa. Comparison of medical and industrial x-ray computed tomography for non-destructive testing. *Case Studies in Nondestructive Testing and Evaluation*, 2016.

[21] Yuval Duchin, Reuben R Shamir, Remi Patriat, Jinyoung Kim, Jerrold L Vitek, Guillermo Sapiro, and Noam Harel. Patient-specific anatomical model for deep brain stimulation based on 7 tesla mri. *PloS one*, 13(8):e0201469, 2018.

[22] Amr Ragab El-Beialy, Mona Salah Fayed, Ahmed Mohammed El-Bialy, and Yehya A Mostafa. Accuracy and reliability of cone-beam computed tomography measurements: Influence of head orientation. *American journal of orthodontics and dentofacial orthopedics*, 140(2):157–165, 2011.

[23] Yang Gao, Borui Li, Wei Wang, Wenyao Xu, Chi Zhou, and Zhanpeng Jin. Watching and safeguarding your 3d printer: Online process monitoring against cyber-physical attacks. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–27, 2018.

[24] Thomas Gregory, Ulrich Hansen, Monica Khanna, Celine Mutchler, Saik Urien, Andrew A Amis, Bernard Augereau, and Roger Emery. A ct scan protocol for the detection of radiographic loosening of the glenoid component after total shoulder arthroplasty, 2014.

[25] Daniel Halperin, Thomas S Heydt-Benjamin, Benjamin Ransford, Shane S Clark, Benessa Defend, Will Morgan, Kevin Fu, Tadayoshi Kohno, and William H Maisel. Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses. In *2008 IEEE Symposium on Security and Privacy (sp 2008)*, pages 129–142. IEEE, 2008.

[26] Vikki Harmonay. Maintaining your ct scanner calibration. <https://info.atlantisworldwide.com/blog/maintaining-your-ct-scanner-calibration>, Jan 2023.

[27] Daniel Ho, Andrew Squelch, and Zhonghua Sun. Modelling of aortic aneurysm and aortic dissection through 3d printing. *Journal of Medical Radiation Sciences*, 64(1):10–17, 2017.

[28] Dietmar W Hutmacher. Scaffolds in tissue engineering bone and cartilage. *Biomaterials*, 21(24):2529–2543, 2000.

[29] Witold Jamróz, Joanna Szafraniec, Mateusz Kurek, and Renata Jachowicz. 3d printing in pharmaceutical and medical applications—recent achievements and challenges. *Pharmaceutical research*, 35(9):176, 2018.

[30] Takashi Kamio, Kamichika Hayashi, Takeshi Onda, Takashi Takaki, Takahiko Shibahara, Takashi Yakushiji, Takeo Shibui, and Hiroshi Kato. Utilizing a low-cost desktop 3d printer to develop a “one-stop 3d printing lab” for oral and maxillofacial surgery and dentistry fields. *3D printing in medicine*, 4(1):6, 2018.

[31] Djim Kanters, Anke de Vries, Henk Boon, Joost Urbach, Arjen Becht, and Homme-Auke Kooistra. Quality assurance in medical 3d-printing. In *World Congress on Medical Physics and Biomedical Engineering 2018: June 3-8, 2018, Prague, Czech Republic (Vol. 1)*, pages 669–674. Springer, 2019.

[32] Tapas Kanungo, David M Mount, Nathan S Netanyahu, Christine D Piatko, Ruth Silverman, and Angela Y Wu. An efficient k-means clustering algorithm: Analysis and implementation. *IEEE transactions on pattern analysis and machine intelligence*, 24(7):881–892, 2002.

[33] Kyobum Kim, Andrew Yeatts, David Dean, and John P Fisher. Stereolithographic bone scaffold design parameters: osteogenic differentiation and signal expression. *Tissue Engineering Part B: Reviews*, 2010.

[34] Robert Lanza, Robert Langer, Joseph P Vacanti, and Anthony Atala. *Principles of tissue engineering*. Academic press, 2020.

[35] Clifford Levi, Joel E Gray, Edwin C McCullough, and Robert R Hattery. The unreliability of ct numbers as absolute values. *American Journal of Roentgenology*, 139(3):443–447, 1982.

[36] Sizhuang Liang, Saman Zonouz, and Raheem Beyah. Hiding my real self! protecting intellectual property in additive manufacturing systems against optical side-channel attacks. In *Proceedings 2022 Network and Distributed System Security Symposium*. Internet Society, 2022.

[37] Daniel A Low, William B Harms, Sasa Mutic, and James A Purdy. A technique for the quantitative evaluation of dose distributions. *Medical physics*, 25(5):656–661, 1998.

[38] Dana Mackenzie. Ultraviolet light fights new virus. *Engineering (Beijing, China)*, 6(8):851, 2020.

[39] Chris Mason and Peter Dunnill. A brief definition of regenerative medicine. 2008.

[40] Materialise. Mimics innovation suite. <https://www.materialise.com/en/medical/mimics-innovation-suite>, June 2017.

[41] Karen Mendelson and Frederick J Schoen. Heart valve tissue engineering: concepts, approaches, progress, and challenges. *Annals of biomedical engineering*, 34(12):1799–1819, 2006.

[42] Henry Miller and Sally Satel. We urgently need new approaches to obtaining organs for transplantation. <https://www.acsh.org/news/2023/02/14/we-urgently-need-new-approaches-obtaining-organs-transplantation-16866>, Feb 2023.

[43] Yisroel Mirsky, Tom Mahler, Ilan Shelef, and Yuval Elovici. CT-GAN: malicious tampering of 3d medical imagery using deep learning. In Nadia Heninger and Patrick Traynor, editors, *28th USENIX Security Symposium*, pages 461–478. USENIX Association, 2019.

[44] Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 2010.

[45] So-Yeon Park, Noorie Choi, Byeong Geol Choi, Dong Myung Lee, and Na Young Jang. Radiological characteristics of materials used in 3-dimensional printing with various infill densities. *Progress in Medical Physics*, 2019.

[46] Sanjana Patrick, N Praveen Birur, Keerthi Gurushanth, A Shubhasini Raghavan, Shubha Gurudath, et al. Comparison of gray values of cone-beam computed tomography with hounsfield units of multislice computed tomography: An in vitro study. *Indian Journal of Dental Research*, 28(1):66, 2017.

[47] Steve Pieper, Michael Halle, and Ron Kikinis. 3d slicer. In *2004 2nd IEEE international symposium on biomedical imaging: nano to macro (IEEE Cat No. 04EX821)*, pages 632–635. IEEE, 2004.

[48] Ben Redwood. Medical 3d printing applications. <https://www.3dhubs.com/knowledge-base/medical-3d-printing-applications/>.

[49] Allied Market Research. 3d printing healthcare market by component, technology, application, and end user: Global opportunity analysis and industry forecast, 2019–2026, Aug 2019.

[50] Antoine Rosset, Luca Spadola, and Osman Ratib. Osirix: an open-source software for navigating in multidimensional dicom images. *Journal of digital imaging*, 17(3):205–216, 2004.

[51] Reihaneh Rostami, Fereshteh S Bashiri, Behrouz Rostami, and Zeyun Yu. A survey on data-driven 3d shape descriptors. In *Computer Graphics Forum*, volume 38, pages 356–393. Wiley Online Library, 2019.

[52] Sanjay Saini, Raju Sharma, Leonard A Levine, Robert T Barmson, Patrick F Jordan, and James H Thrall. Technical cost of ct examinations. *Radiology*, 2001.

[53] N Shahrubudin, P Koshy, J Alipal, MHA Kadir, and TC Lee. Challenges of 3d printing technology for manufacturing biomedical products: A case study of malaysian manufacturing firms. *Heliyon*, 2020.

[54] Jacob C Snyder, Curtis K Stimpson, Karen A Thole, and Dominic J Mongillo. Build direction effects on microchannel tolerance and surface roughness. *Journal of Mechanical Design*, 137(11), 2015.

[55] Jeremy Straub. A combined system for 3D printing cybersecurity. In Kevin G. Harding and Song Zhang, editors, *Dimensional Optical Metrology and Inspection for Practical Applications VI*, volume 10220. International Society for Optics and Photonics, SPIE, 2017.

[56] Johan WH Tangelder and Remco C Veltkamp. A survey of content based 3d shape retrieval methods. *Multimedia tools and applications*, 39:441–471, 2008.

[57] Adam Thompson, Nicola Senin, Claudiu Giusca, and Richard Leach. Topography of selectively laser melted surfaces: a comparison of different measurement methods. *CIRP Annals*, 66(1):543–546, 2017.

[58] Tsung-Hsuan Tsai, N Jeyaprakash, and Che-Hua Yang. Non-destructive evaluations of 3d printed ceramic teeth: Young’s modulus and defect detections. *Ceramics International*, 2020.

[59] Ed Tyson. How to 3d print hips filament – and it’s not just a support material. <https://rigid.ink/blogs/news/how-to-3d-print-hips-filament-and-it-s-not-just-a-support-material>, Oct 2016.

[60] Joseph S Uzarski, Yun Xia, Juan CI Belmonte, and Jason A Wertheim. New strategies in kidney regeneration and tissue engineering. *Current opinion in nephrology and hypertension*, pages 399–405, 2014.

[61] G Harry van Lenthe, Henri Hagenmüller, Marc Bohner, Scott J Hollister, Lorenz Meinel, and Ralph Müller. Non-destructive micro-computed tomography for biological imaging and quantification of scaffold–bone interaction *in vivo*. *Biomaterials*, 28(15):2479–2490, 2007.

[62] Tia Vialva. 3d printed kidneys save life of two-year-old child. <https://3dprintingindustry.com/news/3d-printed-kidneys-save-life-two-year-old-child-133541/>, May 2018.

- [63] Herminso Villarraga-Gómez, Artem Amirkhanov, Christoph Heinzl, and Stuart T Smith. Assessing the effect of sample orientation on dimensional x-ray computed tomography through experimental and simulated data. *Measurement*, 178:109343, 2021.
- [64] Dejan V Vranic and Dietmar Saupe. *3D model retrieval*. PhD thesis, 2004.
- [65] D.V. Vranic and D. Saupe. 3d model retrieval. In *Proc. Spring Conference on Computer Graphics and its Applications (SCCG2000)*, pages 89–93. Comenius University, 2000.
- [66] Mark Yampolskiy, Anthony Skjellum, Michael Kretzschmar, Ruel A. Overfelt, Kenneth R. Sloan, and Alec Yasinsac. Using 3d printers as weapons. *Int. J. Crit. Infrastruct. Prot.*, 14(C):58–71, September 2016.
- [67] Zhiyuan Yu, Zack Kaplan, Qiben Yan, and Ning Zhang. Security and privacy in the emerging cyber-physical world: A survey. *IEEE Communications Surveys & Tutorials*, 23(3):1879–1919, 2021.
- [68] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In *European Conference on Computer Vision*, pages 766–782, 2016.

A Extended Security Analysis

A.1 Proof on 2ϵ -Geometrically Bounded Attacks

The proof of Theorem 1 that XCheck provides 2ϵ -geometrically bounded guarantee is as follows.

Proof. Given a CT scanner with a resolution of ϵ , the distance between any adjacent voxel centroid (v_1, v_2) is bounded by ϵ , i.e. $d(v_1, v_2) \leq \epsilon$. Furthermore, the distance of any voxel within the volume of model is also bounded by ϵ , otherwise, a new voxel should be created to describe the volume.

Given the assumption above, we can prove Theorem 1 by contradiction. Without loss of generality, we assume that an adversary \mathcal{A} modifies the PSD, which causes a deviation of magnitude 2ϵ and remains undetected. Then $\exists v_o \in M_d, p_n \in V_n, d(v_o, p_n) > 2\epsilon$ by definition 1, where V_n is the new volume. Since all voxels must be bounded by ϵ , a new voxel has to be added to describe the new volumetric region, V_n . Because the shortest distance of p_n to any existing voxel is greater than 2ϵ (i.e. $\min[d(p_n, v)] > 2\epsilon, v \in M_d$), and the distance of all voxels in V_n must be bounded to ϵ , therefore new voxels must be created to describe the volume. As a result, for the new voxel v_n , it must satisfy $d(v_n, p_n) < \epsilon$. Together with the constraint that $d(v_o, p_n) > 2\epsilon$, we can deduce that $d(v_n, p_n) > \epsilon$, since within the plane that is formed

by the triangle of (v_o, v_n, p_n) , the sum of two sides must be greater than the other. On the other hand, by the definition of being undetected by ϵ -geometrically bounded comparison, $\max[Dev(V_n, M_d)] < \epsilon$, and because V_n is a subset of M_c , therefore $\forall p \in V_n, \max[D(p, M_d)] < \epsilon$. This is a contradiction to the fact that $d(v_n, p_n) > \epsilon$, $v_n \in V_n, p_n \in V_n$. Therefore it is impossible for an attacker to craft an attack that causes a volume deviation exceed 2ϵ yet remain undetected by XCheck, under the assumption that CT reconstructed the model M_n of device with resolution of ϵ . The same process can be used to prove modifications that lead to removal of voxels. \square

A.2 Detecting Stealthy Adversarial Attacks on Geometry

We assume the adversary has access to both the algorithm and parameters of XCheck and therefore they can launch white-box attacks. While we have proved that geometry attacks above 2ϵ can always be detected by XCheck, when the attack is below that threshold, the detection becomes probabilistic.

For a single voxel on the compared model, if the attacker intends to evade voxel analysis while maximizing its deviation, the attacker’s goal can be expressed as:

$$\max M, \text{s.t. } |M + e_1 + e_2| \leq T_{vox}, e_1 \sim f_1, e_2 \sim f_2 \quad (4)$$

where T_{vox} represents the threshold for voxel analysis, f_1 and f_2 are the probability density functions of e_1 and e_2 . The distribution of the attack magnitude can be expressed as:

$$f_{vox}(M) = \iint f_1(e_1) f_2(e_2) de_1 de_2, \text{s.t. } |M + e_1 + e_2| \leq T_{vox} \quad (5)$$

As such, the attacker may estimate a probability density function of the attack magnitude, where a larger magnitude will likely to cause more severe damage, while a higher density means the corresponding magnitude will be more likely to evade detection. The attacker may make a trade-off to select the proper attack magnitude. Assume the attacker selects the attack magnitude to be within a range $[m, n]$, then the possibility of evading detection for each point can be estimated as $p = \int_m^n f_{vox}(M) dM$. As the voxel analysis checks every voxel and only passes if all the deviations are within threshold T_{vox} , the possibility of a malicious PSD to evade is $P_{vox} = \prod_{i=1}^n p_i$. Note that the increase in false positive rate due to this approach is mitigated by the gamma analysis. In a benign print, even if the printed model may have several voxels that deviate from the original design, the exceed magnitude is very small.

In order to evade ray-based analysis, the attacker may estimate the volumetric difference v_{ray} in a similar formulation as Eq. 5. Since ray-based analysis accumulates and restricts deviations in arbitrary regions, the attack space for evading detection is further compressed. As a result, our analysis based on different types of geometric invariants significantly limits the attacker’s ability to modify PSDs without being detected, thus lowering the risk of cyber-physical attacks on patients.

Table 3: Gamma analysis for detecting geometry and material attacks

Base Model	Attacked Feature	Attack Extent	Gamma Analysis					Run Time (sec)	Attack Detection	
			S	D	V	M	G			
Bone Screw	Length	0.2cm	1.28	1.10	1	1	2.20	235.8	✓	
		0.4cm	1.57	1.16	1	1	2.41	236.2	✓	
		0.8cm	2.61	1.40	1	1	3.28	235.3	✓	
		1.2cm	3.55	1.54	1	1	4.12	234.7	✓	
	Thread distance	0.06cm	1	1.21	1	1	2.11	230.8	✓	
		0.15cm	1	1.41	1	1	2.23	232.4	✓	
		0.3cm	1	1.47	1	1	2.27	233.7	✓	
		0.5cm	1	1.57	1	1	2.34	236.4	✓	
	Minor diameter	0.03cm	1	1	1	1	2	236.6		
		0.06cm	1	1.10	1	1	2.05	237.1	✓	
		0.18cm	1	1.24	1	1	2.13	236.2	✓	
		0.28cm	1	1.25	1	1	2.14	234.7	✓	
		0.54cm	1	1.38	1	1	2.21	232.7	✓	
	Thread section	1	1	1.17	1	1	2.09	238.5	✓	
	Material replacement	LCD-C	1	1	1	6.51	6.74	238.8	✓	
		LCD-E	1	1	1	8.46	8.64	236.9	✓	
		LCD-N	1	1	1	5.72	5.98	238.2	✓	
		LCD-SG	1	1	1	8.62	8.79	235.6	✓	
		LCD-MG	1	1	1	5.94	6.19	239.4	✓	
	Benign #1	-	1	1	1	1	2	233.1	(TN)	
	Benign #2	-	1	1	1	1	2	235.3	(TN)	
Bone Scaffold	Solid region	0.025cm	1	1.03	1	1	2.02	278.1		
		0.05cm	1	1.43	11.93	1	12.10	269.3	✓	
		0.075cm	1	1.51	13.98	1	14.13	271.4	✓	
		0.1cm	1	1.89	20.4	1	20.54	278.6	✓	
	Hollow region	0.025cm	1	1.29	3.78	1	4.24	266.8	✓	
		0.05cm	1	1.31	7.62	1	7.86	278.2	✓	
		0.075cm	1	1.40	14.63	1	14.76	269.6	✓	
		0.1cm	1	1.46	25.71	1	25.79	277.5	✓	
	Porous region	0.025cm	1	1.38	1.23	1	2.33	279.9	✓	
		0.05cm	1	1.44	3.17	1	3.76	287.8	✓	
		0.075cm	1	1.51	6.33	1	6.66	278.3	✓	
		0.1cm	1	1.63	8.25	1	8.53	277.5	✓	
	Condyle angle	5 degree	1	1.3	1	1	2.17	270.1	✓	
		10 degree	1	1.6	1	1	2.36	278.4	✓	
		20 degree	1	1.68	1	1	2.41	267.3	✓	
		30 degree	1	2.08	1	1	2.71	288.1	✓	
	Benign #3	-	1	1	1	1	2	268.8	(TN)	
Lung-on-a-chip	Material replacement	LCD-C	1	1	1	6.68	6.90	287.4	✓	
		LCD-N	1	1	1	3.36	3.78	268.7	✓	
		LCD-MG	1	1	1	5.45	5.72	280.3	✓	
	Benign #4	-	1	1.08	1	1	2.04	277.6	✓(FP)	
	Benign #5	-	1	1	1	1	2	274.4	(TN)	
	Benign #6	-	1	1.13	1	1	2.07	280.2	✓(FP)	
Dental Guide	Material replacement	LCD-G	1	1	1	5.38	5.65	278.2	✓	
		Benign #7	-	1	1	1	1	2	276.1	(TN)
		Benign #8	-	1	1	1	1	2	276.5	(TN)