

Genuinity Signatures: Designing Signatures for Verifying 3D Object Genuinity

Daniel G. Aliaga Mikhail J. Atallah

Department of Computer Science at Purdue University

ABSTRACT

3D computer graphics models and digitally-controlled manufacturing have come together to enable the design, visualization, simulation, and automated creation of complex 3D objects. In our work, we propose and implement a framework for designing computer graphics objects and digitally manufacturing them such that no adversary can make imitations or counterfeit copies of the physical object, even if the adversary has a large number of original copies of the object, knowledge of the original object design, and has manufacturing precision that is comparable to or superior to that of the legitimate creator of the object. Our approach is to design and embed a signature on the surface of the object which acts as a certificate of genuinity of the object. The signature is detectable by a signature-reading device, based on methods in computer graphics and computer vision, which contains some of the secret information that was used when marking the physical object. Further, the compromise of a signature-reading device by an adversary who is able to extract all its secrets, does not enable the adversary to create counterfeit objects that fool other readers, thereby still enabling reliable copy detection. We implemented a prototype of our scheme end-to-end, including the production of the physical object and the genuinity-testing device.

Categories and Subject Descriptors (according to ACM CCS): I.3 [Computer Graphics], I.3.3 [Picture/Image Generation], I.3.7 [Three-dimensional Graphics and Realism], I.4.1 [Digitization and Image Capture].

1. Introduction

Our work provides algorithms for encoding into a digital 3D object information that enables determining genuinity of a physical object after its automated manufacturing. In today's technological world, many physical objects are manufactured using 3D computer graphics models and digitally-controlled devices (e.g., milling machines, 3D printers, and robotic arms). The manufactured objects can range from inexpensive steel screws to costly and carefully designed parts, for example, for engines and for medical instruments. Our algorithms provide a way to encode a unique signature into a computer-designed and fabricated object and a way to decode the signature in order to verify object genuinity.

The digitization of the 3D design and manufacturing process of objects enables creating intricate structures that may have required a significant investment of time, infrastructure, and personnel. A growing worry amongst designers, manufacturers, and buyers of digitally-built parts is knowing whether they have a genuine part fabricated by the advertised company. While for simple objects a visual inspection might be sufficient to detect irregularities, this is not the case for more complex objects or for verifying material composition, density, and product quality. Thus, imitations and "knock-offs" may float around the black market and fall into unexpected hands. The replicas are designed to

unknown specifications, which may lead to catastrophic problems (e.g., failures in an engine). A recent manufacturing industry report [<http://mems.org>] called counterfeiting "the crime of the century". Even in 1997, the cost of counterfeit parts to the global automotive industry was already \$12 billion, and the numbers are up sharply since then.

The challenge is to encode information into the object that cannot be reproduced by an adversary and to do so without depending on the legitimate manufacturer having technology superior to the adversary or on security-through-obscurety. The high-resolution etching of logos or serial numbers is an option but can be imitated by an adversary that develops equal or better technological ability than the legitimate manufacturer -- a fact that can be true without the legitimate manufacturer knowing so. An instance of security-through-obscurety would be in assuming the digital model itself or the location of unique markings on the object is kept secret. However, if this information leaks out, security is compromised and the adversary is able to make copies indistinguishable from genuine objects.

Our main observation is that upon copying any physical object, the adversary cannot obtain the exact same combination of manufacturing errors. For digital data, perfect copies can be obtained; thus digital watermarking is a known process for encoding ownership data into a digital object. However, our problem is quite different from digital watermarking and investigates an under-explored area of research. For physical objects, copying is always an imperfect

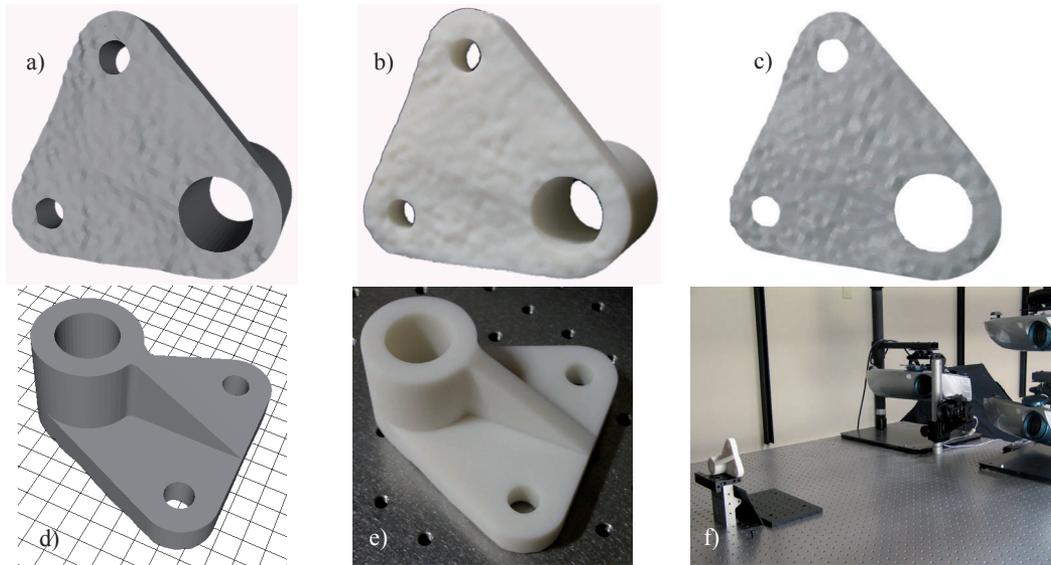


Figure 1: Genuinity Signatures. We address the problem of detecting whether a physical object is genuine or is a counterfeit. (a, d) Using our algorithm, the user designs a unique “genuinity signature” and encodes it into a subset of a 3D synthetic model. (b, e) The physical object is built using automated digital manufacturing. (c, f) An automated system determines if the object is genuine or is an imitation (even if manufactured using higher accuracy than the genuine).

and noisy process. We interpret the change of manufacturing error as a change in the “signature” of the object. Thus, we design a method to inject a small amount of additional and carefully designed manufacturing error into a subset of the object’s surface (that is inconsequential to its intended use). During a verification process, we reliably discover unexpected error and determine object genuinity.

Our approach creates a genuinity signature for an object by using a 3D synthetic model of the object, a set of accuracy values for the manufacturing and verification processes, and an arbitrarily-chosen object serial number. The serial number and a private “master key” are used, for example, in a cryptographically secure one-way hash (e.g., SHA-2) to generate apparently random numbers for defining the genuinity signature. The signature is composed of a set of random yet smoothly varying surface displacements that, under user control, is applied to the signature’s footprint on the object (i.e., a subset of or the entirety of the object’s surface). The secret key assists in generating random numbers that determine both where and how the signature is embedded, but reverse engineering the key is hard because there are an exponential number of possibilities. Moreover, even if the adversary was willing to exhaustively try all possibilities, the adversary has no clear “success criterion” to know when a correct guess has been made.

The displacements have been carefully designed so that upon either *re-instancing* (i.e., making copies from scratch using the original digital model) or *replication* (i.e., making copies by digitally acquiring and then re-manufacturing the object) the signature is distorted. To verify genuinity of an object, our process acquires a high-resolution surface model of the signature’s footprint on the object and uses statistical tools to determine if the signature is still present. Further, our signature is also designed to include some redundancy

so that very precise alignment of the captured footprint and the synthetic signature is not necessary. Figure 1 shows the process for an exemplary object.

Our algorithm does not rely upon the legitimate manufacturer having technology superior to the adversary nor on storing secret information, except for the aforementioned master key. In fact, in our work we assume that the entire encoding/decoding algorithm and the digital model itself are public knowledge. The master key can be specific to an object or group of objects. Thus, even if the master key becomes public knowledge, only that object or small group of objects is compromised. Moreover, for verification only a partial-key that decodes some of a particular signature is needed. If this key becomes public, an adversary cannot copy the entire affected object or objects.

In more detail, let’s consider the case of a genuine object that can be manufactured with variance $m_g > 0$ and an adversary able to produce a similar object with a manufacturing variance $m_a > 0$. Given an operational tolerance τ , we can assume $m_g \leq \tau$ and $m_a \leq \tau$; but, we do not know if $m_g \leq m_a$ or vice versa. In our work, subsets of the object’s surface are intentionally displaced by up to $(\tau - \sqrt{m_g})$. If the object has been re-instanced or replicated, then the observed displacements and their distribution is different than expected because either the desired displacements were unknown or a variance m_a appeared unexpectedly. Even if m_a is smaller than m_g , the number of displacements can be made larger so as to reliably detect a smaller m_a .

To demonstrate our method, we design several objects, encode a signature, either digitally manufacture the object using 3D printing facilities or using simulation, and then capture and verify genuinity. Our main contributions are

- an approach enabling the detection of physical counterfeit objects created in today's world of digital design and manufacturing,
- a specific encoding method to "write" a unique signature into a physical 3D object without affecting the object's functionality, in the sense of not exceeding an operational error tolerance, and
- a general decoding method to "read" a signature from a physical 3D object such that if the object has been copied, even with higher accuracy than the manufacturer's technology, the signature cannot be extracted. The reading process does not have to reconstruct the entire object, is not highly sensitive to alignment, and does not need a priori calibration.

2. Related Work

Before the rampant digitization of the manufacturing process, the authenticity of an object could be determined by the presence of a unique set of marks at a particular (possibly inconspicuous) location, e.g., a special etching, a serial number, etc. The digitization of manufacturing made many of these schemes obsolete. In the past the adversary could produce the same markings on the physical object while not necessarily creating objects to the same specifications of quality, shape, and material compositions as the original. Today's counterfeiters can often have the same manufacturing technology, or better, and without the original manufacturer even knowing it. Other approaches in use today consist of attaching to the physical object an electronic device (e.g., RFID [RVW*07]) or a smartcard coupled with a holographic-watermarking scheme. While our method can be used without such attached electronics (AEs), an adversary could make copies of the AE device and attach each such copy to a superior-quality counterfeit copy. Using our method in coordination with AEs prevents such attacks because the AE would no longer match the object to which it is attached.

A related effort is copy-evident technology and physical watermarking for secure paper documents (e.g., [Vol*96]). These approaches either use a printing resolution greater than a typical paper copier to embed special markings onto a document or embed special light absorption and reflection material into the paper itself making a copied document very evident. By visual inspection or by using a magnifying glass the markings can be visually verified.

Watermarking has been extended to the digital domain (e.g., [CGE*07, WLD*07]) and is concerned either with robustness (e.g., [LDD*07, PYC*03]) or with fragility (e.g., [CHL*01, YY*99]) for ensuring the integrity of images and other digital objects. These approaches have been designed for digital data which can be created, read, and replicated with zero error. In contrast, for physical objects there is an error involved at each of the aforementioned steps. Although the second group of watermarking methods could be adapted to our targeted problem, the schemes would fail our requirement of selective robustness in the face of our own manufacturing errors, and the inaccuracies of our verification tools. Traditional uses of fragile watermarking have not had a counterfeiting adversary in mind. Their model of the adversary was of one who wishes to modify the object without causing the disappearance of the fragile mark – it was

```

CreateGenuinitySignature (Key K)
  Select signature footprint on object.
  Define per-pixel displacements  $\mathbf{d}_i$ 
    and variances  $\mathbf{v}_i$ .
  Divide pixels into contiguous patches  $\mathbf{A}_i$ .
  Sort patches by their variances.
  Define non-adjacent but equal-variance
    groups  $\mathbf{B}_k$ .
  Subdivide model into small triangles and
    perturb footprint using displacement map.

PerformGenuinityTest (Key K)
  Acquire 3D model of signature footprint.
  Align with synthetic model of displaced
    surface geometry.
  Compute value of the test statistic  $\mathbf{W}$ 
    to determine genuinity.
  
```

Figure 2: Pseudocode Summary of Method.

really anti-tamper. But our adversary has no interest in tampering. Instead, they seek to produce a faithful reproduction of the object, mark, and all.

3. Genuinity Signatures

A genuinity signature is composed of a large set of samples organized by a two-level stratification which has been designed to robustly detect physical imitations created via re-instancing or replication (Figure 2). Each sample in the signature stores both a surface displacement value and a surface variance value. The former helps to create a signature with a unique configuration and the latter helps to detect unwanted replication attacks. The organization into strata provides additional structure that improves the sensitivity for detecting attacks yet can still maintain the appearance of a random signature. All sample values and stratification are determined using a secure key-based pseudorandom number generator and randomly parameterized Perlin noise [Per*02]. Perlin noise is employed in order to generate smoothly changing displacement and variance values. The smoothness and stratification also has the side-effect of relaxing the alignment requirement between the synthetic and captured signature thus making verification easier.

The large set of samples is beneficial to enabling an abundant number of unique signatures and to yielding the precision needed to guard against an adversary with technology superior to that of the legitimate manufacturer. We chose to use samples consisting of local displacements, rather than global alterations (e.g., of volume, surface area, or ratios of the sizes of different parts of the object), because verification can be performed without reconstructing the entire object. Furthermore, the size of the signature's footprint on the object is proportional to the desired robustness against attacks. Recovering a surface fragment of the object visible from a single viewpoint is often sufficient.

Our method can be applied to any object as long as some visible surface area can be manufactured with our specified displacements. The absolute value of the errors is not important. Instead, the effectiveness of genuinity detection depends on the number and grouping of the displacements and on the ratio between the accuracies of the legitimate's and adversary's technology. Thus, our method can be used with both low-end and high-end manufacturing technology.

3.1 Error Model

While not all fabrication processes can be used to generate a surface fragment with detailed displacements, we restrict fabrication of the subset of the object intended for the signature to a suitable method. These processes are assumed to be imperfect and are modeled by an unbiased normal error distribution, and of a standard deviation that is pre-specified or determined experimentally.

We define the genuine object's manufacturing variance m_g and verification variance v_g as the square of a chosen multiple of the standard deviation of the corresponding process. Specifically, $m_g = (k\sigma_m)^2$ and $v_g = (k\sigma_v)^2$ for some value of $k > 0$, σ_m is the standard deviation of the legitimate manufacturer's technology, σ_v is the standard deviation of the verification technology (e.g., the 3D acquisition method). Thus, for $k = 3$ approximately 99% of the object will be manufactured within $\sqrt{m_g}$ of ideal. Further, by using variances (and not standard deviations), we can state that the total variance of the system is $m_g + v_g$. Also, the difference between the operational tolerance τ and $\sqrt{m_g}$ provides the freedom to introduce surface displacements.

The corresponding values for the adversary are defined as fractions of the legitimate manufacturer's: $m_a = (1/\beta)m_g$ and $v_a = (1/\beta)v_g$. We succinctly refer to the adversary having technology β times better than the legitimate creator.

3.2 Signature Samples

To obtain the displacement d_i and variance v_i for each sample i , we use Perlin noise ρ parameterized by the secret key K (Figure 3a-b). The key is used to generate parameters for Perlin noise and to generate the random numbers within the procedure. Thus, sample $s_i = \{(d_i, v_i)\}$ and

$$\begin{aligned} d_i &= 2m\rho(K) - m \\ v_i &= 2r\rho(K) - r \end{aligned} \quad (1)$$

where $\rho(K)$ generates Perlin noise in the range $[0,1]$, m is the sample displacement magnitude, and r is the desired sample variance value. For $Z = XY$ and $i \in [1, Z]$, both d_i and v_i are stored in 2D arrays of size $X \times Y$.

Displacement and variance values are collectively used to compute the amount by which to displace a surface fragment along the direction of the surface normals. To keep the surface area of the signature within operational tolerance,

the signature sample must satisfy

$$m + \sqrt{r} \leq (\tau - \sqrt{m_g}). \quad (2)$$

Uniqueness of the signature is accomplished by having the displacement image be of non-trivial size (e.g., a 100x100 resolution displacement image already yields a huge number of possible signatures). The smoothness of the displacements and variances is implicitly obtained via Perlin noise. We control the range of smoothness of the noise by limiting the cloud density and cloud coverage values.

3.3 Signature Strata

We organize the samples s_i into a two-level stratification: a spatial stratum and a variance stratum (Figure 3c). The spatial stratum defines patches A_j where each patch contains samples spatially adjacent to each other. The variance stratum B_k defines groups of patches such that the variance value of each group is nearly constant but the variance changes from group to group. Each stratum contains a reference to all signature samples. Signature S can be compactly written as

$$\begin{aligned} S &= [\{A_j: j \in [1, N]\}, \{B_k: k \in [1, M]\}] \\ A_j &= \{s_{i_1j}, s_{i_2j}, s_{i_3j}, \dots\} \\ B_k &= \{A_{j_1k}, A_{j_2k}, A_{j_3k}, \dots\} \end{aligned} \quad (3)$$

containing N patches of the form A_j in the spatial stratum and M groups of the form B_k in the variance stratum. All the elements of a particular A_j are spatially adjacent and all the elements of a particular B_k have (almost) the same variance value stored in all contained samples.

Spatial Stratum

Re-instanting is guarded against by the use of the key-based displacement amounts. However, the spatial stratum is helpful to reduce the need to precisely align the synthetic version of the signature with the acquired signature during verification and thus account for tangential errors in measuring and in manufacturing. Since both the displacement and variance values are generated using a smooth noise function, the aforementioned patches provide some slack which implies that the synthetic version of the signature does not need to be perfectly aligned with the captured object. For example, a bad alignment or unexpected tangential distortion (with respect to the reading device) might cause a small subset of the samples of one patch to be compared against

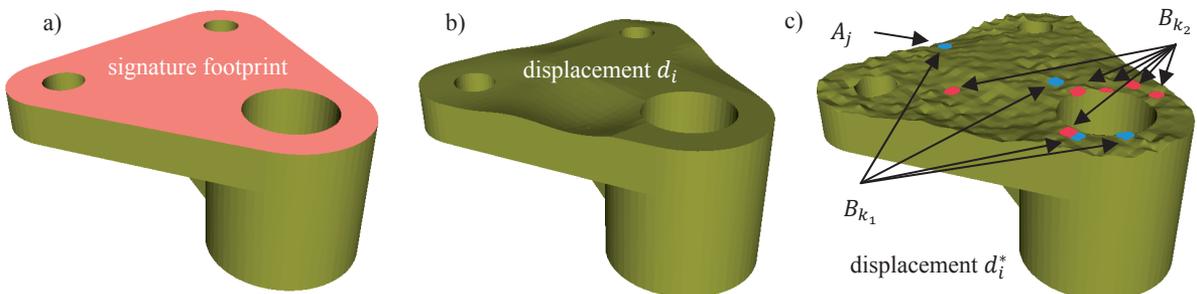


Figure 3: Genuinity Signatures. a) User selects a signature footprint. b) Random but smooth per-pixel displacements d_i are computed. c) A smoothly-varying additional variance v_i per pixel is calculated, yielding a total displacement of d_i^* . Adjacent samples are joined into patches A_j and spatially disjoint but of equal variance patches into groups B_k .

samples of an adjacent patch. While this is not desirable, the use of smooth noise makes it such that the adjacent patch has a different but similar value.

Variance Stratum

The variance stratum is useful to make verification more sensitive to a change of the amount of variance as a result of an attack. Changing the variance amongst the groups exacerbates the need to know the correct way to group the patches (and samples). The likelihood of randomly grouping patches and samples to yield the same particular set of variances is extremely small. For example, if a replication was performed and all groups in the signature had no additional variance added then all the copied object groups would exhibit homogeneity of variance. While the variance is not exactly equal to that of the legitimate creator's manufacturing and verification processes (because of the additional variance introduced by the adversary's replication), the homogeneity of all the variances makes it more difficult to detect an overall change of the amount of variance. In contrast, by providing each group with a different additional variance, homogeneity of variance is not obtainable unless the correct grouping is known.

3.4 Creation Algorithm

We use a three-phase algorithm to generate the genuinity signature and to apply it to an arbitrary object.

- First, signature image size $X \times Y$ and signature strength values m and r are chosen (subject to inequality (2)) and then equations (1) produce displacement values and variance values per image sample.
- Second, the patches A_j of the spatial stratum are created. They are determined by joining a contiguous block of $w \times h$ samples in the signature image into a single unit. This yields patches consisting of different displacement values and, as will be computed in the next step, a constant variance value. When $w = h = 1$, each sample maps to a surface fragment on the object consisting of many adjacent vertices. This yields patches with a smooth change in the displacement value (e.g., using a bilinear or bicubic interpolation) and a constant variance value. This setup relaxes even more the need for accurate alignment during verification. Since the number of samples can be made to be sufficient, this variant is usually more advantageous.
- Third, the groups B_k of the variance stratum are selected. To form the groups, all patches are sorted by their variance values and then divided into M groups each of approximately \hat{Z}/M samples, where $\hat{Z} = XY/(wh)$. To ensure spatially-near variances are similar, we compute a group's variance \bar{v}_k as the average of the variances of all contained samples (rather than randomly picking a group variance value), namely

$$\bar{v}_k = \frac{1}{\sum_{j \in B_k} |A_j|} \sum_{j \in B_k} \sum_{i \in A_j} v_i. \quad (4)$$

To map the signature to the object, the user selects the signature footprint as a subset of the object's surface and then maps the signature image to the footprint. In our prototype, selection is done via a simple interactive process of clicking-and-selecting on triangles. The selected triangles

are recursively subdivided such that no triangle edge is larger than $\sqrt{m_g}$. The new triangulation is at the resolution of the manufacturing process. As previously mentioned, the mapping between signature image pixels and object vertices does not need to be one-to-one. Rather, the signature image resolution can be less than the resolution of the object triangulation in order to yield patches of samples of smoothly-changing displacement but constant variance value. While several image-to-object mappings are possible, our prototype uses orthographic projection. Further, since the signature footprint is not necessarily rectangular, clipping may occur when mapping the rectangular signature image to an object; nevertheless, the loss of samples is not problematic because the signature can be made suitably larger so that the actual number of mapped samples is sufficient.

The total per-sample displacement, stored in the image, is

$$d_i^* = d_i + N(0, \sqrt{\bar{v}_{\kappa(i)}}) \quad (5)$$

where $N(\mu, \sigma)$ returns a random value with a normal distribution of mean μ and standard deviation σ , $\kappa(i)$ returns the index of the group containing s_i . Given P points in the signature footprint, object point x_p ($p \in [1, P]$) in general corresponds to fractional positions on the signature image. The total displacement d_p^* from x_p and along the surface normal n_p is calculated by a bilinear equation of the surrounding displacement values in the signature, e.g.

$$d_p^* = (1-s)d_{i_1}^* + sd_{i_2}^* + (1-t)d_{i_3}^* + td_{i_4}^* \quad (6)$$

where s and t are the fractional positions of x_p on the displacement image and $(d_{i_1}^*, d_{i_2}^*, d_{i_3}^*, d_{i_4}^*)$ are the surrounding displacement values. The new model is composed of points

$$x_p^* = x_p + n_p d_p^* \quad (7).$$

4. Genuinity Testing

To verify genuinity, we need a digital copy of the original unmodified object model, the key (or the relevant part thereof), and the physical 3D object. Either the manufacturer or a client can perform the verification. However, while the digital object model and signature creation and verification algorithms can be public, the key must be kept secure. In addition, we desire genuinity testing to be fast, accurate, and easy-to-use. These characteristics permit a widespread process that can be encapsulated into a simple "reader".

It is worth noting that because of the potentially large number of groups, not all of the groups B_k are needed for genuinity testing. Instead, a random subset of them suffices. Thus, we can design the signature to have an excessive number of groups such that only a subset of them is needed. This redundancy yields two abilities: 1) the entire signature does not have to be reconstructed for genuinity testing – reducing the burden for the verification system, and 2) genuinity testing can be provided with only a "partial key", which we define as a subset of the bits of the master key and decodes only a subset of the groups and samples. This yields added security because if the partial-key becomes public, the object can still not be fully imitated because the rest of the genuinity signature is not known.

4.1 Acquisition

To verify the genuinity of the signature, we need to capture the signature footprint. Since the displacements are usually subtle, high-resolution acquisition is needed. However, we note that the macro-structure of the area to capture is usually quite simple (e.g. planar or a smooth curved section). Thus, one suitable option is to use a photogeometric capture approach. Namely, capture the general geometric structure and then refine the details using a photometric-based reconstruction. Photometric-based approaches are good at obtaining fine details but not so good for accurately obtaining the global structure. Photogeometric approaches have been successfully used before to obtain reconstructions accurate up to a fraction of a millimeter. For example, the Digital Hammurabi project [KSD*04] performed a 3D scanning of cuneiform tablets at 0.025mm accuracy (1/1000 inches = 0.025mm). [NRD*05] used a geometric and a photometric setup to obtain high accuracies (between 0.01mm and 0.1 mm). The recent examination of the Mona Lisa done by NRC (Canada) used a custom 3D color-scanner of 0.06x0.06x0.01 mm resolution [BGM*07]. [AX*08] created a self-calibrating approach that used projectors as both light sources and virtual cameras yielding multi-viewpoint photogeometric reconstructions sampled at 0.05mm. These example works show that reconstructions of sufficient quality are possible.

Photogeometric Capture

In our prototype, we created a “signature reader” similar to the aforementioned photogeometric capture approaches. Our particular system is based on [AX*08] in the sense of being self-calibrating and combining both geometric and photometric data, however our method is single-viewpoint (like [NRD*05]). The surface fragment is placed approximately perpendicular to the optical axis of the camera and thus can be reconstructed at the full resolution of the photometric method (i.e., that of the camera), which in our case is about 10x times greater than that of the projectors (Figure 1f). We refer the reader to [AX*08] for more details.

Signature Alignment

The presence of sample patches and the smoothness of the displacement value relax the need for very accurate alignment between the physical object and the synthetic version of the object (perturbed by d_p^*). As is typically the case with single-viewpoint capture methods observing an object head-on, the dominating error is in the distance measurements along the camera ray direction (and not errors in the tangential directions). Thus, the local smoothness of the signature displacements helps to diminish the effect of not precisely aligning the synthetic object with the physical object – a small tangential displacement results in only small distance measurement changes. Moreover, this relaxation also compensates for small manufacturing errors tangential to the displacement direction.

In our prototype, the user interactively places the captured surface fragment over its synthetic equivalent. Then, iterative-closest-point (ICP) is used to further refine the alignment [RL*01]. The fragment is resampled so that each point \tilde{x}_p on the fragment is the closest possible point to the cor-

responding point x_p^* on the synthetic model. Only the surface fragment corresponding to the signature needs to be processed. Also, if the surface and signature are very different, ICP might fail. This is not problematic since that most likely indicates object genuinity should be rejected.

4.2 Verification

The objective of the verification process is to ensure that the signature is sufficiently intact so as to conclude the object is genuine. We use statistical tests to yield a nearly continuous probability value indicating object genuinity. First, we describe the use of a simple similarity of means test. This test is straightforward and easy to implement but is not always sensitive enough for replication attacks. Second, we describe a more sophisticated test that is able to robustly and correctly reject a re-instanced or replicated object.

Similarity of Means Testing

Re-instancing can be detected with the use of a key-based generation of displacements and a similarity of means statistical test. Since the key is kept secure, it is extremely unlikely that the adversary can recreate all the displacement values of the signature. A re-instanced object will most likely have a different signature and thus a simple similarity of means test is sufficient. Using the synthetic model with the signature displacements applied, each synthetic point x_p^* is subtracted from the corresponding reconstructed object point \tilde{x}_p . The mean of the magnitude of these difference vectors should be zero for all reconstructed points. A z-test or Student’s t-test can be used to check for a mean with a statistically significant difference from zero.

Group-based Sampling and Median-based Testing

In our verification tests, we use the stratification of the samples and median-based testing to determine homogeneity of variance [Lom*07]. A replicated object will contain an unexpected additional variance within the captured groups B_k that was introduced by the adversary. Re-instancing will also cause the captured variance within each group B_k to be different than the expected variance. Intuitively, we desire two properties: 1) the variance of all samples in a group to be the same and 2) the variance of each group to be \bar{v}_k plus $m_g + v_g$. We express this by using a modified version of a Brown-Forsythe test which checks for homogeneity of variance (i.e., homoscedasticity) using medians and variable transformation. The use of medians, as opposed to means, is known to be more robust to samples with an error distribution straying from perfect normalcy (e.g., a simple F-test is rarely useful unless the distribution is perfectly normal).

We modify the Brown-Forsythe test so that it tests for homogeneity of variance and for equality of that variance to a desired variance. The signature samples of each group B_k mapped onto a genuine object are composed of a displacement d_i , a group variance \bar{v}_k , and an additional ($m_g + v_g$) variance introduced by the manufacturing and verification processes. Once we account for the first two components, on average the group should exhibit values with a variance equal only to the third component. The magnitude of \bar{v}_k changes from group to group and its value is \bar{v}_k only if we group the points correctly. Thus, obtaining a leftover va-

riance of $(m_g + v_g)$ is only the case if the proper organization of the strata is known, the sample displacements are known, and the only additional variance is $(m_g + v_g)$. Moreover, since we have a large number of groups, we can robustly test if the variance is in fact $(m_g + v_g)$.

If the test statistic exceeds a chosen threshold value, then it indicates rejection (i.e., object is re-instanced or replicated). For this test, the transformed measurement variable for group k and element $p \in B_k$ is $y_{kp} = \tilde{x}_{kp} - x_{kp}^*$ and $\bar{y}_k = \text{median}\{y_{kp}; p \in B_k\}$. Then, we define the terms

$z_{kp} = |y_{kp} - \bar{y}_k|$, absolute difference with group median,

$z_k = \frac{1}{S_k} \sum_{p \in B_k} z_{kp}$, mean of group k , and

$z = (m_g + v_g)$, desired overall variance

where S_k is the number of elements in each group. These terms are used in the test statistic

$$W = \frac{(P-M) \sum_k S_k (z_k - z)^2}{(M-1) \sum_k \sum_{p \in B_k} (z_{kp} - z_k)^2} \quad (8)$$

which if it exceeds $F(\alpha, M-1, P-M)$ means the homogeneity of variance fails with $M-1$ and $P-M$ degrees of freedom and at a significance level α (F is a function that returns the upper critical value of an F-distribution). For convenience, we swap the numerator and denominator of equation (8) in order to make $W \geq 1$ at all times.

Mutual Verification

One potentially problematic situation is if both the legitimate manufacturer and the adversary attempt to include a genuinity signature in the same object; who will be recognized as the genuine manufacturer? This situation is, however, not problematic with our approach. If manufacturer B acquires and copies the object from manufacturer A and produces a new object with their own genuinity signature, then the new object will fail the genuinity test for manufacturer A and succeed for that of manufacturer B . This is the desired and appropriate behavior.

5. Results and Discussion

We have used our algorithms to encode and verify genuinity signatures on several types of objects. First, to encode a signature our software system reads in a description of a 3D model. A digital key is provided and used to generate the unique genuinity signature. Second, the newly created model is either fabricated using a stereolithography process or provided to a simulator. Our simulator enables us to thoroughly explore the parameter space. Third, a reconstruction of an object is obtained and checked for the presence of the genuinity signature. Our prototype verification system uses a Canon Rebel XTi 10 MP camera and three Optoma EP910 DLP projectors of 1400x1050 pixels each.

In the following, we present results including several graphs that analyze the behavior of the signature and its verification in a given situation. The horizontal axis in all graphs represents the number of groups used when designing the signature. The vertical axis represents the value of the test statistic of equation (8). Each individual curve corresponds to either a particular size of the signature displacement image or to a particular verification scenario.

Each datapoint is the averaged result of 10 repeated simulations of creating a signature image, perturbing it according to the legitimate manufacturer's and adversary's relative accuracies for manufacturing and verification, and performing a genuinity test. While arbitrarily more groups can be used by increasing the total number of samples, our graphs focus on the minimum number of samples and on the minimum (or maximum) number of groups needed for a given sample size. Also, since our methods scale to any accuracy level, we arbitrarily fix $\tau = 1$ and adjust other parameters.

The simplest interpretation of the numerical value of the test statistic (y-value of the graphs) is that a value near 1 implies success of the genuinity test. Given the total number of samples, the number of groups, and the test statistic value, we can compute the statistical significance level at which the genuinity test will fail. This value varies with each datapoint in the graphs and can be computed using statistical tables or software equivalents. While we do report it for some examples, usually a test statistic value near 2 or greater strongly implies the object is not genuine.

We have analyzed in simulation the behavior of the signature and its verification for various parameter values. While we have performed the simulations for a range of values, in these representative results we assume that the operational tolerance is about twice the manufacturing accuracy and there is a 5:1 ratio between m and \sqrt{r} ; more precisely $\tau = 1$, $m_g = v_g = 0.5$, $m = 0.25$ and $r = 0.0025$ – these numbers guarantee the operational tolerance is not exceeded by the signature displacements and manufacturing process. Figure 4a shows the effectiveness of our method in recognizing a truly genuine object using several signature image sizes. Once a sufficient number of groups are used (i.e., ≥ 60), genuinity testing robustly succeeds in keeping the test statistic near 1. Figure 4b shows the ability of our genuinity test in discovering a re-instanced object. Using a 200x200 signature image or larger, yields sufficient samples to robustly detect the re-instanced object. The effectiveness of the test increases as the number of groups reduces because each group is assigned more samples resulting in a stronger change of variance. However, the desire to have fewer groups with more samples runs counter to the need for more groups that is needed for verifying truly genuine objects. For a signature image size of 200x200 about 60 to 100 groups yields a compromise between the two scenarios. More groups can be used for larger signature image sizes.

The use of the signature variance helps make the test more robust to attacks. Figures 4b-c report the test statistic using a ratio of 5:1 and 1:5 between m and \sqrt{r} . For all cases, a variance difference between groups improves the robustness of the tests but yields a less smooth signature that is difficult to fabricate and to maintain undamaged.

Figures 4d-f show the behavior of our approach during replication attempts for an adversary with $\beta = 1$, $\beta = 10$ ("10 times better technology") and $\beta = 100$ ("100 times better technology"). In these tests, the relative values of m_g , v_g , s_d , and s_p are not important – the value for β , as well as the image size and group count, dominate the behavior. For $\beta = 1$, a 50x50 displacement image is sufficient. For $\beta = 10$, at least a 250x250 image is needed to robustly

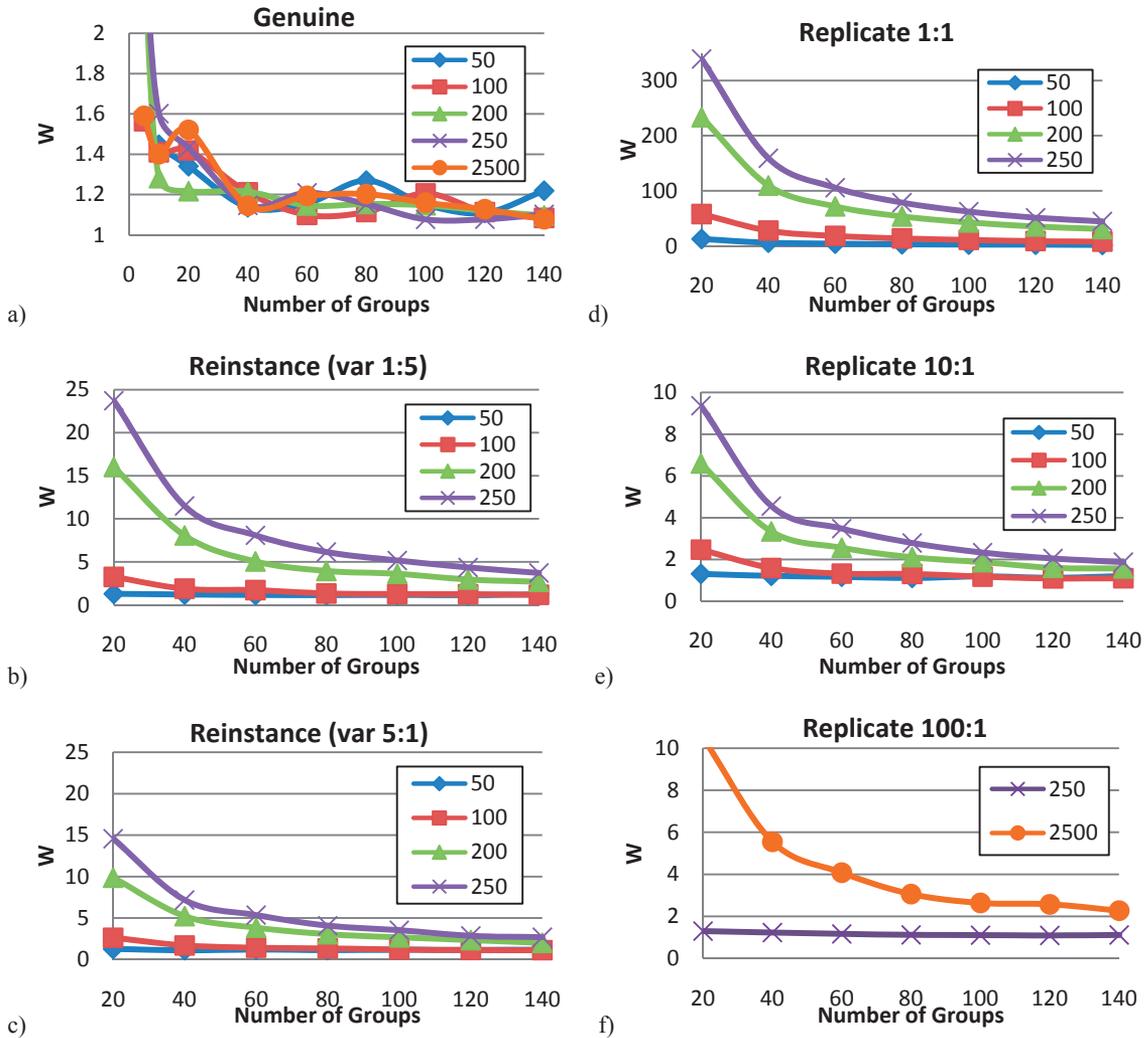


Figure 4: Genuinity Testing Analysis. We show the behavior of our genuinity test for a) genuine objects, b-c) for re-instanced objects, and d-f) for replicated objects. The two reinstancing examples also show the benefit of having a variance component in the signature; b) show a ratio of 1:5 between m and \sqrt{r} and c) show a ratio of 5:1. The replication attacks use an adversary with technology d) 1x, e) 10x, or f) 100x better than the legitimate’s manufacturer.

reject such a replicated object. For $\beta = 100$, a larger signature displacement image of 2500x2500 is needed. In all cases, using up to about 140 groups is possible. Hence, even if the adversary has manufacturing and acquisition accuracy 100 times superior to the legitimate manufacturer, the replicated object can still be robustly identified.

In Figure 5, we summarize the behavior of our genuinity signature and its verification in several “attacks” by an adversary. Figure 5a shows the behavior of the genuinity test when using a signature displacement image of 250x250 samples and $\beta = 10$. Figure 5b shows give a similar report but for $\beta = 100$ and using a signature displacement image of 2500x2500 (note: the test-value for the re-instance attack is very large, for convenience we divide it by 1000). In all cases, our approach clearly discerns the genuine object.

Figure 1 (on the first page) shows a synthetic object, a fabricated object, and an acquired fragment used for genuinity

testing of a CAD object. Figure 1a and 1d show the 3D computer model. The signature is visible on the bottom side. Figure 1b and 1e are *photographs* of the fabricated physical object. Figure 1c is the acquired surface fragment of the signature. Since our genuinity detection is designed for fragility upon copying (i.e., success is hard), our real-world experiments test accurately recovering the signature at the expected levels of error. The effects of re-instancing and replicating an object can be accomplished, respectively, by comparing the same captured object to a different signature and by changing the variances of the signature.

The signature we design and verify uses a displacement image of 35x35 samples and 100 groups. The square displacement image is mapped onto the object with only simple scaling and orthographic projection, thus only about 2/3’s of the samples are actually used. The object measures 102x100x45 mm. The accuracy of the stereolithography fabrication device available to use can be set between 0.1

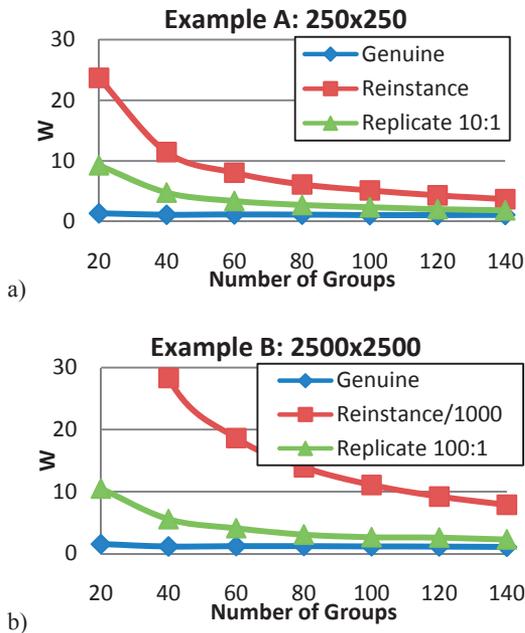


Figure 5: Example Attacks. a) Attempt of copying using $\beta = 10$ and a 250x250 signature displacement image. b) Another set of attacks using $\beta = 100$ and a 2500x2500 signature displacement image.

mm and 0.2 mm (i.e., about 5/1000th to 8/1000th of an inch) -- we used 0.1 mm. Our verification system is estimated to have an accuracy of 0.5 mm; in this example, 99% of the acquired fragment was at most 0.39 mm from ideal (thus $v_g = 0.5$ is reasonable). Our manufacturing and verification variances are taken as $m_g = 0.15^2$ and $v_g = 0.5^2$. For this prototype, we satisfy $\tau = 2$ mm by using $m = 1.0$ mm and $r = 0.75$. The averaged test statistic value for genuine objects, re-instanced objects, and replicated objects (using $\beta=10$) is, respectively, 1.26, 3.57, and 2.16. In this signature, the upper critical F-value for the involved degrees of freedom and at a 5% significance level is about 1.26. Thus, the truly genuine object passes the test and the re-instanced and replicated object are correctly rejected with only a 5% chance of getting the classification wrong.

Figure 6 shows the relative effect of the signature parameters m and r . In this example, we assume we have 1 mm of “space” for the signature displacements. m provides the general shape of the signature (Figure 6a: high m and low r). r provides variability between the groups (Figure 6b: low m and high r). High variability between the groups

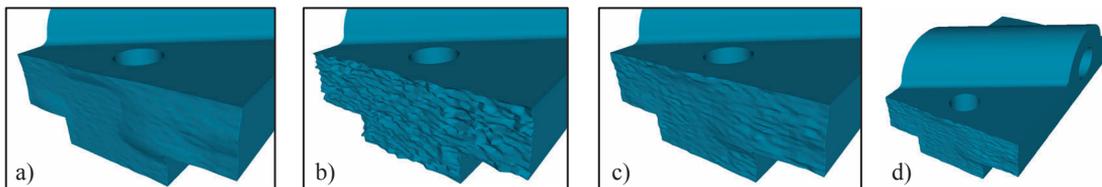


Figure 6: Signature Design. We show a CAD object with the signature at both end caps for $\tau = 1$ mm. a) Signature with relatively high m and low r , b) signature with low m and high r and c-d) a balanced combination.

yields the most robust test, but the signature itself might be hard to manufacture and/or fragile. Our choice is to use a balance of the two parameters (Figures 6c-6d). In this small footprint signature, the significance level is about 10% (e.g., the test may be wrong 10% of the time – this number can be improved by further tuning the parameters).

In Figure 7, we demonstrate the flexibility of our method by applying it to a previously scanned object rather than a CAD object. The process for our method is identical and the signature in fact looks very much like existing noise.

Figure 8 shows another experimental example. This object was fabricated with similar parameters as that in Figure 1 but using $\tau = 2$ and $m_g = 0.2^2$ mm (“standard stereolithography quality” of our supplier). The test statistic values for genuine objects, re-instanced objects, and replicated objects (using $\beta=10$) is, respectively, 1.23, 2.75, and 1.63. The acquired object’s was at most 0.47 mm from ideal.

We further validate our method by showing the effect on the signature for replicated objects. Using the objects from Figures 1 and 8, we re-manufacture the objects from the acquired version (containing the signature). In Figure 9, we show the difference (using a jet colormap) between corresponded points of the digital model with signature (i.e., no errors) and scanned surface fragments. The difference represents the error introduced by manufacturing and verification -- it should be relatively smooth and small. Figures 9a and 9c use the acquired signatures of the genuine objects where test statistic $W=1.16$ and $W=1.22$, respectively. Figures 9b and 9d use the scanned signatures of the replicated objects where $W=2.86$ and $W=2.43$, respectively -- thus the objects are detected to be replicas. Further, Figures 9a and 9b are created using the same high-resolution manufacturing technology; replica Figure 9d is created using high-resolution manufacturing technology while the corresponding genuine object is created using the lower-quality standard-resolution process. Nevertheless, the replication process introduces additional error that can be detected. (Note: the straight-line artifacts of 9a and 9b are due to the capture methods; the red-bands on the sides of 9c and 9d are because that surface region is not part of the signature.)

6. Discussion and Future Work

We have presented a novel approach to encode data for detecting object genuinity into the surface of physical objects. Our problem is quite *different* from digital watermarking where perfect copies can be made. In the physical world, copying is a noisy process that we exploit to detect replication attempts, re-instancing attempts, or both. Our approach encodes a signature into a digital 3D object which after

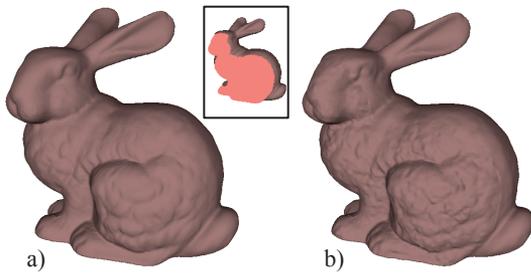


Figure 7: Scanned Object. a) Scanned object, no signature. b) Object with signature added (middle figure highlights the signature footprint).

manufacturing is used to determine genuinity. Our work provides a unique blend of computer graphics, information security, and advanced manufacturing.

Our method does have some limitations and aspects that need further refinement. First, our approach requires “space” to place the signature. Such space might not exist on all types of objects. Second, our error model is based on an assumption of a near-normal error distribution. While we explicitly use techniques to support deviations from normality, a severe departure will likely confuse verification. Third, our technique requires approximate knowledge of the variance of the manufacturing process and verification process. A significant over (or under) estimation of these quantities will also distort the verification procedure. While the variance can be determined experimentally, we look to more sophisticated statistical processing to support using conservative estimates. To provide more robustness and to support a wider range of object types and manufacturing processes, a promising extension is to use other feature spaces (e.g., color) and their combinations.

7. References

- [AX*08] ALIAGA D., XU Y., “Photogeometric Structured Light: A Self-Calibrating and Multi-Viewpoint Framework for Accurate 3D Modeling”, *IEEE Computer Vision and Pattern Recognition*, 2008.
- [BGM*07] BORGEAT L., GODIN G., MASSICOTTE P., POIRIER G., BLAIS F., BERARDIN J.A., “Visualizing and Analyzing the Mona Lisa”, *IEEE Computer Graphics & App.*, 27(6), 60-68, 2007.
- [CGE*07] CORSINI M., GELASCA E.D., EBRAHIMI T., BARNI M., “Watermarked 3-D Mesh Quality Assessment,” *IEEE Trans. on Multimedia*, 9(2), 247-256, 2007.
- [CHL*01] CHEN M., HE Y., LAGENDIJK R.L., “Error detection by fragile watermarking”, *Picture Coding Symp.*, 287-290, 2001.
- [KSD*03] KUMAR S., SNYDER D., DUNCAN D., COHEN J., COOPER J., “Digital Preservation of Ancient Cuneiform Tablets Using 3D-Scanning”, *3-D Digital Imaging and Modeling*, 2003.
- [LDD*07] LAVOUE G., DENIS F., DUPONT F., “Subdivision Surface Watermarking,” *Computers & Graphics*, 32(3), 480-492, 2007.
- [Lom*07] LOMAX R., An Introduction to Statistical Concepts, *Routledge*, 472 pages, 2007.
- [NRD*05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHY R., “Efficiently Combining Positions and Normals for Precise 3D Geometry”, *ACM SIGGRAPH, TOG* 24(3), 2005.
- [Per*02] PERLIN K., “Improving Noise”, *ACM Transactions on Graphics*, 21(3), 681-682.

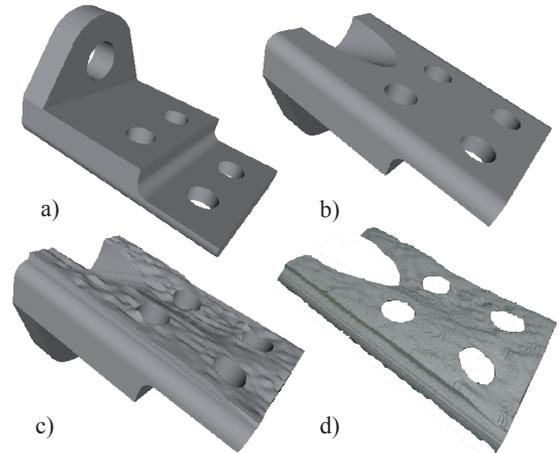


Figure 8: Additional Example. a) Top view of original 3D synthetic model. b) Bottom view of original model. c) View of model with encoded signature for $\tau = 2$ mm. d) Acquired signature footprint.

- [PYC*03] PENG X., YUB L., CAIB L., “Digital watermarking in 3D space with a virtual-optics imaging modality”, *Optics Communications*, 226(1-6), 155-165, 2003.
- [RL*01] RUSINKIEWICZ S., LEVOY M., “Efficient Variants of the ICP Algorithm”, *3D Digital Imaging and Modeling*, 145-152, 2001.
- [RVW*07] RIDA A., VYAS R., WU T., LI R., TENTZERIS M., “Development and Implementation of Novel UHF Paper-Based RFID Designs for Anti-counterfeiting and Security Applications”, *IEEE Int'l Workshop on Anti-counterfeiting, Security, & Identification*, 52-56, 2007.
- [Vol*96] VOLPE H.R., “Printing method and copy-evident secure document”, *US Patent 5487567*, January, 1996.
- [WLD*07] WANG K., LAVOUE G., DENIS F., BASKURT A., “3D Meshes Watermarking: Review and Attack-Centric Investigation”, *Int'l Workshop on Information Hiding*, 50-64, 2007.
- [YY*99] YEO B.-L., YEUNG M.M., “Watermarking 3D Objects for Verification”, *IEEE Computer Graphics & Applications*, 19(1), 36-45, 1999.

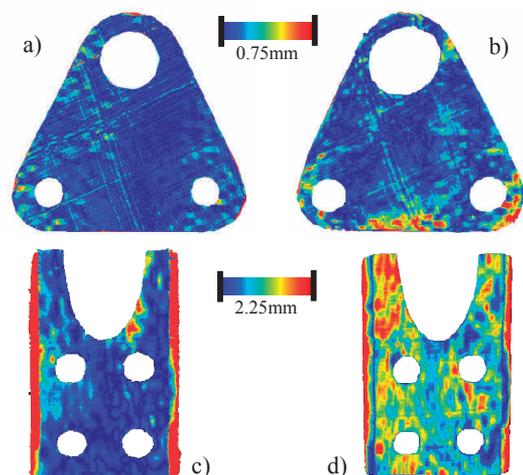


Figure 9: Replication Examples. Differences are shown between digital model with signature and surface fragments (using jet-colormap). (a, c) Captured signatures of genuine objects; (b, d) replicated objects.