

ANTI-CROPPING BLIND RESYNCHRONIZATION FOR 3D WATERMARKING

Xavier Rolland-Nevière, Gwenaël Doërr

Pierre Alliez

Technicolor R&D France

Inria Sophia Antipolis - Méditerranée

ABSTRACT

Radial-based 3D watermarking alters the distances between the center of mass of the 3D mesh and its vertices. These watermarking systems are inherently sensitive to cropping. To address this limitation, this paper introduces a complementary blind resynchronization module to transmit critical synchronization information to the watermark decoder. Spherical patterns formed by several secret landmark vertices are embedded alongside the payload and blindly retrieved by the decoder, thereby conveying the synchronization information needed. Experimental results showcase significant improvement against cropping, while preserving performances against valumetric attacks thanks to a control parameter that automatically switches between alternate resynchronization modes.

Index Terms— 3D landmarks, self-embedding, robustness.

1. INTRODUCTION

Several 3D watermarking techniques have been proposed and they range from using complex multi-resolution or spectral transforms on meshes, to straightforward alterations of the vertex locations to embed the watermark payload [1]. When the watermark decoder still has access to the original cover mesh, it can rely on advances in 3D shape matching [2] to perform a non-blind resynchronization prior to retrieving the payload from a potentially damaged input. This being said, in many application use cases, watermark decoding is required to be blind and thus raises technical challenges, e.g., to survive cropping attacks [3].

Typical countermeasures against watermark desynchronizing attacks include: (i) using an auxiliary pilot sequence for realignment purposes, (ii) using a watermark carrier oblivious to the loss of synchronization, and (iii) exploiting resynchronization patterns formed by mesh features [4, Chapter 9]. In the context of 3D, cropping-invariant carriers have been reported to be sensitive to valumetric attacks such as noise addition [5]. Today, most proposals rely on the third countermeasure. The features of the mesh guide a segmentation procedure and the payload is repeatedly embedded in every segment. It is likely that some segments are left untouched by cropping, thus allowing for a payload extraction. Defining stable features (umbilical points [6], prongs [7, 8]) and cropping-invariant segmentations remain enduring challenges. Moreover, with this strategy, the watermark designer cannot control the number and location of the feature points; the robustness is highly content dependent. Finally, feature points are public, which provides opportunity for security attacks.

In this paper, we propose an alternate blind resynchronization approach that is akin to a pilot sequence. We illustrate its relevance for a family of so-called *radial-based 3D watermarking systems* [9, 10, 11]. The latter embed the watermark payload by modulating the distribution of radial distances ρ_i between the vertices of the mesh and its center of mass \mathbf{g} . Successful payload extraction

for such systems requires three content-dependent pieces of information, namely \mathbf{g} and the bounds $M = \max(\rho_i)$ and $m = \min(\rho_i)$ of the radial distances, which are often altered after cropping. In 3D watermarking, off-line transmission of this critical information to the detector, a.k.a. semi-blind watermarking, is ineffective to deal with combinations of cropping and rigid transforms or scaling.

In a nutshell, the proposed resynchronization approach relies on the introduction of secret landmarks on the surface mesh in such a way that their geometric arrangement encodes the desired critical information. In Section 2, we first define what is a secret landmark and then detail how such secret landmarks can be introduced in a 3D mesh and subsequently blindly recovered. Using this baseline component, Section 3 then describes how to place these secret landmarks in a specific geometric arrangement, e.g., onto a sphere, to convey the desired critical information. Moreover, we discuss how to leverage on some landmark-dependent indicators to automatically switch between alternate resynchronization modes depending on the attack context. The proposed resynchronization module is finally integrated into a reference radial-based watermark algorithm [11] in Section 4. Reported benchmarking results clearly indicate superior performances against cropping while preserving performances against valumetric attacks. Section 5 provides concluding remarks as well as directions for future work.

2. LANDMARK VERTICES

2.1. Definitions

Given a field $\mathbf{f} : \mathcal{M} \rightarrow \mathbb{R}^n$ on a surface mesh, a landmark vertex is defined as a vertex v whose *signature* $\mathbf{f}(v) = [f_1(v) \dots f_n(v)]^T$ is close to an element $\mathbf{q} = [q_1 \dots q_n]^T$ of a lattice $\mathcal{Q} \subset \mathbb{R}^n$. In other words, the set of landmarks \mathcal{L} of a mesh corresponds to the vertices such that $\|\mathbf{f}(v) - \mathbf{q}(v)\| < \delta$ holds true, where $\mathbf{q}(v)$ is the element of \mathcal{Q} closest to $\mathbf{f}(v)$ and δ is a threshold distance. In the case of orthogonal lattices, $q_i = k\Delta_i + \epsilon_i$ ($k \in \mathbb{Z}$), where Δ_i is a quantization step and ϵ_i is a secret offset. $q_i(v)$ is computed as:

$$q_i(v) = \lfloor (f_i(v) - \epsilon_i)\Delta_i^{-1} + 0.5 \rfloor \Delta_i + \epsilon_i. \quad (1)$$

Introducing a landmark at a given vertex v amounts to modifying the local geometry around the vertex so that its signature $\mathbf{f}(v)$ falls within the positive detection subspace, i.e., the union of spheres of radius δ and centered at all \mathbf{q} in \mathcal{Q} . In general, this is difficult to achieve, especially when the signature depends on a large neighborhood around v . In this study, this limitation motivates the use of a compact signature based on fitting the 2-ring vertex neighborhood of v to a paraboloid surface [12]. The neighbor vertices are first mapped to a local basis, centered at v , and whose axes are given via Principal Component Analysis (PCA). The parametric model $z(x, y) = ax^2 + bxy + cy^2 + dx + ey + f$ is then fit using linear

least squares, and the signature is taken as:

$$\mathbf{f}(v) = \left(\frac{a^2 + 2b^2}{|a + c|^2}, \frac{c^2 + 2b^2}{|a + c|^2} \right). \quad (2)$$

Using the 2-ring keeps the neighborhood compact while providing a sufficient number of vertices to avoid ill-defined cases during least squares fitting. Moreover, this 2-dimensional signature is invariant to rigid transforms, uniform scaling, as well as cropping, provided that it does not impact the 2-ring around v .

2.2. Introduction and Blind Recovery of Landmarks

As mentioned above, turning an arbitrary vertex into a landmark is a challenging problem, especially because the mapping between the vertices and the signature space depends on the intermediary content-dependent PCA. For instance, it implies that the landmark creation cannot reduce to projecting the vertices of the local neighborhood onto a paraboloid patch having the target parameter values. We thus relocate the neighboring vertices via active-set optimization (`fmincon` [13]), so as to reach the positive detection subspace while minimizing the distortion.

The unknowns are the coordinates of the neighboring vertices and the cost function to minimize is the squared error metric, i.e., the sum of squared vertex displacements. The landmark characteristic is expressed by a non-linear constraint set to $[|f_1(v) - q_1(v)|, |f_2(v) - q_2(v)|]^T$ if $\|\mathbf{f}(v) - \mathbf{q}(v)\| \geq \delta$ and $\mathbf{0}$ otherwise, i.e., when the vertex signature is inside the positive detection region. The solver then minimizes the geometric distortion while keeping this constraint null. For this optimization procedure we could rely upon alternate mesh distortion metrics that are more perceptually tuned [14]. However, these metrics incur significant computational overhead and are unfit for highly localized modifications.

On the receiver side, the detector computes the signature $\mathbf{f}(v)$ for each vertex of the mesh and compares its distance $d(v)$ to the nearest element in \mathcal{Q} to a threshold τ . Inserted landmarks should all have $d(v) \leq \delta$ in the absence of any alteration. As a result, setting $\tau = \delta$ blindly identifies landmark vertices. However, false positives may also be found, and their occurrence can commonly be controlled by (i) tightening the lattice constraint δ ; or (ii) increasing the spatial support of the signature $\mathbf{f}(v)$.

2.3. Experimental Assessment

For evaluation, all quantization steps Δ_i are set to 5×10^{-2} . 50 landmark vertices with non-overlapping 2-ring neighborhoods are arbitrarily placed on the mesh *Bunny*, before being blindly detected. This experiment is repeated for several values of $\alpha = 2\delta/\Delta$ (normalized δ) and, for each setup, the true and false positive rates (TPR/FPR) are recorded when varying the decision threshold τ . Figure 1 depicts these experimental measurements using dotted Receiver Operating Characteristic (ROC) curves. While the FPR may look acceptable at first, one should keep in mind that meshes are made of a very large number of vertices, e.g., approximately 3.5×10^4 vertices for the *Bunny*. As a result, a FPR as low as 10^{-3} nearly yields as many erroneous landmarks as correct ones. In particular, the FPR for the operating point $\tau = \delta$ depicted by a square marker is too high for practical use, even with very small positive detection regions using $\alpha = 0.02$ (dotted green).

To tackle this issue, the proposed landmark creation procedure is applied iteratively on increasingly wider neighborhoods to enlarge

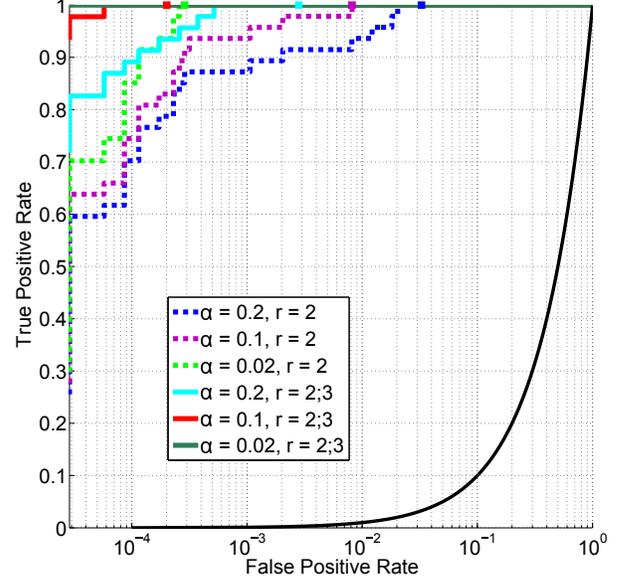


Fig. 1. Classification performances of blind landmark detection, with (solid lines) and without (dotted lines) the nested signature, on *Bunny*, for several normalized lattice constraints α . Random classification is depicted by the black curve; a squared marker indicates the operating point $\tau = \delta$.

the dimension of the signature and thus reduce the likelihood for vertices to be detected as landmarks. To avoid interferences between the nested signatures, neighbor vertices that have been modified are subsequently left untouched. In this study, the dimension of the signature is increased from 2 to 4, thanks to an increase from 2 rings to 3 rings.

The aforementioned process is first applied to the 2-ring. Then, the signature of the 3-ring is estimated and moved to the positive detection region by only modifying the 3-ring vertices that are not in the 2-ring. For simplicity, the same signature definition (Eq. (2)) and the same parameter values are used for both nested neighborhoods. For classification, $d(v) = \sqrt{(d_2(v))^2 + d_3(v)^2}/2$ is compared to a threshold τ , where $d_2(v)$ (resp. $d_3(v)$) is the distance to the lattice \mathcal{Q} computed for the 2-ring (resp. 3-ring) signature. The corresponding ROC curves are depicted with solid lines in Figure 1. For strong closeness constraints ($\alpha = 0.02$), false positives are no longer observed, as indicated by the lack of dark green marker. This setting is chosen as default in the remainder of this paper.

Empirically, the induced local distortion remains negligible, e.g., when measured with the local MSDM [15]. Another observation is that the likelihood for the minimization to fail increases when using smaller values for δ . For instance, using the recommended settings, the nested procedures statistically yields a couple of failures. Nonetheless, this failure rate remains small enough to be manageable in the second part of the resynchronization framework.

3. RESYNCHRONIZATION BASED ON LANDMARKS

As secret landmarks can be arbitrarily introduced on the mesh, the objective is to locate them at specific configurations to transmit critical information for resynchronization, namely \mathbf{g} , m , and M in the

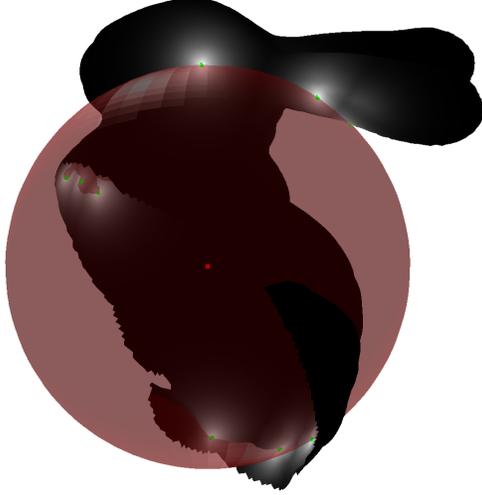


Fig. 2. Blind recovery of the center of mass (red dot) of the *Bunny*. The landmarks inserted at embedding (indicated by white spots on the mesh) are used to fit a sphere whose center coincides by construction with the original center of mass, which cannot be estimated from the cropped mesh (40% cropping ratio).

radial-based 3D watermarking systems. To begin with, only the center of mass \mathbf{g} is transmitted by arranging secret landmarks onto a secret sphere. We further explain how to exploit the *quality* of the identified landmarks to automatically switch between alternate resynchronization strategies depending on the attacking context. Finally, two secret landmarks configurations are combined to retrieve all the critical resynchronization information on the receiver side.

3.1. Transmitting the Center of Mass

The proposed strategy is to arrange landmarks in a specific geometric configuration that can be blindly estimated on the receiver side and whose geometric characteristics encode the desired information. For instance, to transmit the center of mass, landmarks are placed on a sphere centered at \mathbf{g} . Upon detection, (i) landmark vertices are recovered, (ii) a robust sphere fitting is performed, and (iii) the center of the estimated sphere is used instead of \mathbf{g} . This procedure is expected to provide robustness against cropping as long as a sufficient number of landmarks are recovered to accurately estimate the sphere, as depicted in Figure 2.

First, the set $\mathcal{I}(r, \epsilon)$ of vertices whose radial distance is between $r - \epsilon$ and $r + \epsilon$ is identified. r is a predefined radius and $\epsilon > 0$ a control parameter that guarantees that $\mathcal{I}(r, \epsilon)$ contains enough vertices. In practice, ϵ is set to 0.1% of $M - m$. These vertices are close to a sphere $\mathcal{S}(\mathbf{g}, r)$ centered at \mathbf{g} and they are candidates to be turned into landmarks. The next step consists in identifying a subset $\mathcal{L} \subset \mathcal{I}(r, \epsilon)$ of n_L non-coplanar vertices with independent neighborhoods that will be used to form the landmark-based resynchronization pattern. This is an independent set problem that can be tackled with a greedy sequential selection procedure: vertices in $\mathcal{I}(r, \epsilon)$ are added one by one, starting with the ones having the lowest number of neighbor vertices also belonging to the neighborhood of another vertex in $\mathcal{I}(r, \epsilon)$ and vertices whose addition would yield a neighborhood overlap in \mathcal{L} are discarded. This strategy empirically spreads the selected vertices over the surface of the mesh (see

Figure 2), which decreases the likelihood that a cropping attack completely removes the resynchronization pattern. All vertices in \mathcal{L} are projected onto $\mathcal{S}(\mathbf{g}, r)$ to guarantee that \mathbf{g} coincides with the center of the spherical pattern. The landmark creation procedure is finally applied, without altering the position of the vertices in \mathcal{L} .

In practice, this procedure may fail to produce n_L non-coplanar landmark vertices for a number of reasons: there are not enough vertices with non-overlapping neighborhoods in $\mathcal{I}(r, \epsilon)$, the vertices in \mathcal{L} are co-planar, the landmark creation fails, etc. This issue is addressed by scanning $[m, M]$ to find a radius r that provides the desired number of landmarks.

3.2. Robust and Adaptive Blind Registration after Attacks

On the receiver side, the detector identifies the landmark as described in Section 2. In general, the estimated set of landmarks $\hat{\mathcal{L}}$ contains false positive and misses some of the correct ones. As a result, using least-squares sphere fitting to estimate the resynchronization pattern is unreliable. To deal with outliers, a RANSAC [16] is used to estimate the resynchronization sphere. Its estimated center $\hat{\mathbf{g}}_{\hat{\mathcal{L}}}$ replaces the center of mass $\hat{\mathbf{g}}$, which is potentially damaged by cropping, in the watermark detection [11].

The threshold τ provides an indirect means to control the trade-off between the landmark robustness and the accuracy of the estimated resynchronization information. Increasing τ widens the landmark detection region and increases the likelihood of recovering landmarks after valumetric attacks such as noise addition. It also increases the FPR but this is compensated by RANSAC. However, the stability of $\hat{\mathbf{g}}$ is notably higher than the stability of $\hat{\mathbf{g}}_{\hat{\mathcal{L}}}$ in case of valumetric-only attacks. In this context, using the proposed resynchronization approach may be counter-productive, which is a rather common consideration for countermeasures against cropping.

The quality of the landmarks can then serve as an indicator to switch to a fall-back estimation strategy to estimate the center of mass. By setting $\tau = \delta$, the landmark selection is rather strict. In case of cropping-only attacks, landmarks that are still present in the attacked mesh are recovered and the RANSAC sphere estimation is accurate. In contrast, in case of valumetric-only attacks, most landmarks move out of the detection region and they are not recovered. The RANSAC estimation then indicates a poor fitting and the system falls back to estimating the center of mass from the damaged mesh. The same strategy is applied when the RANSAC estimation fails. This strategy allows to tackle both attacking scenarios, although combined valumetric and cropping attacks remain a challenge.

3.3. Conveying Additional Resynchronization Information

To transmit the remaining critical information for resynchronization, namely the bounds of the histogram m and M , a straightforward approach consists in using two distinct spherical synchronization patterns $\mathcal{S}(\mathbf{g}, r_1)$ and $\mathcal{S}(\mathbf{g}, r_2)$ and to exploit the two radii r_1 and r_2 to encode m and M . First, two distinct types of landmarks need to be defined to correctly associate landmarks with their resynchronization patterns. This is achieved by relying upon two distinct lattices \mathcal{Q}_1 and \mathcal{Q}_2 in a manner similar to binary quantization index modulation (QIM [17]). Next, the radii are set as: $[r_1 \ r_2]^T = \mathbf{L}[m \ M]^T$, where \mathbf{L} is a full-rank mixing matrix. In this study, the first row of \mathbf{L} is set to $[0.5 \ 0.5]$ and the second is set to $[0.6 \ 0.4]$. The sequential constructions of the set of candidate landmarks \mathcal{L}_1 and \mathcal{L}_2 are then modified to guarantee that there is no neighborhood overlap between both sets of vertices. On the receiver side, the information

derived from the spheres estimated using RANSAC ($\hat{\mathbf{g}}_{\hat{\mathcal{L}}}$, $\hat{m}_{\hat{\mathcal{L}}}$, and $\hat{M}_{\hat{\mathcal{L}}}$, where $\hat{\mathcal{L}} = \hat{\mathcal{L}}_1 \cup \hat{\mathcal{L}}_2$) is used for watermark decoding rather than the quantities directly derived from the mesh. When a single one of the two resynchronization patterns cannot be recovered, the center of the only remaining sphere $\hat{\mathbf{g}}_{\hat{\mathcal{L}}}$ is used, and \hat{M} and \hat{m} are directly computed from the mesh.

4. EXPERIMENTAL BENCHMARK

For performance evaluation, we consider a reference radial-based 3D watermarking algorithm that formulates the payload embedding process as an optimization problem solved with quadratic programming [11]. The only modification is that the landmarks, as well as their neighbor vertices, are left untouched to avoid damaging the inserted resynchronization pattern(s). Four different variants of the baseline system are then studied:

1. without any resynchronization;
2. with a single resynchronization pattern encoding \mathbf{g} ;
3. with two resynchronization patterns encoding \mathbf{g} , m , and M ;
4. with the ground truth resynchronization information.

The last variant provides a theoretical upper bound on the performances that could be achieved. For benchmarking, a database containing 12 meshes is used. Their number of vertices ranges from 5k to 100k. The embedding strength parameter is adjusted to calibrate the embedding distortion (MSDM [15]) at 0.15. The meshes are watermarked 5 times with random 64-bits payloads and each watermarked mesh is randomly cropped 7 times to produce as many attacked meshes that are input to the decoder. Figure 3 depicts the average Bit Error Rate (BER) recorded for the resulting 180 meshes for various cropping ratios.

In line with previously reported benchmarking results, the baseline system is extremely sensitive to cropping due to the loss of critical information (\mathbf{g} , m , and M). In contrast, the solid green curve using ground truth side information indicates how much gain could be achieved with an efficient resynchronization module. The BER could remain below 5% even for strong cropping attacks. Adding a single resynchronization pattern to recover the original center of mass already provides significant performances improvement. As a matter of fact, the BER actually alternates between 0% and 50% depending on whether the bounds of the histogram m and M are recovered or not. The stronger the cropping attack, the more likely these bounds are lost. As a result, incorporating the second resynchronization pattern to convey m and M further improves robustness. The dotted red curve remains very close to the lower bound provided by the ground truth, until it rockets above 10% deletion. For such strong cropping attacks, many landmarks are lost. As a result, the resynchronization approach is less likely to recover both patterns and the bounds \hat{m} and \hat{M} used for watermark decoding are then corrupted. This explains why the curves for the two variants with resynchronization patterns nearly coincide for large cropping ratios.

For completeness, the evaluation also includes testing robustness against a wide range of attacks such as rigid transforms, scaling, noise addition, smoothing, quantization, simplification, refinement or triangle soups [11]. Albeit not depicted due to the lack of space, the two variants that incorporate the proposed resynchronization module showcase similar performances as the baseline system. This validates that setting $\tau = \delta$ effectively prevents the synchronization module from jeopardizing the established large robust-

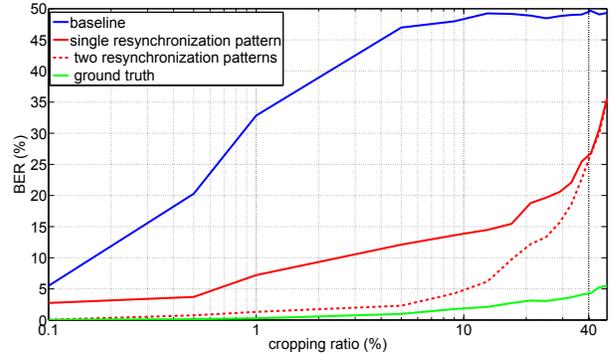


Fig. 3. Robustness performances against cropping for several variants of the baseline radial-based 3D watermarking system and increasing cropping ratios (log-scale).

ness of the baseline watermarking system against valumetric-only attacks.

5. CONCLUSION

We presented a resynchronization approach that deals with cropping attacks in 3D watermarking. The main idea relies upon inserting landmarks on a mesh in a specific geometric configuration, so that critical resynchronization information can be computed from the landmarks recovered blindly from the receiver side. Conceptually, this approach is similar to pilot sequences which have been proposed in the past, except that the sequence itself is content-dependent and used to convey resynchronization information. The proposed approach is evaluated for a reference radial-based 3D watermarking system [11] and the reported experimental results showcase a significant boost of robustness against cropping. Contrary to other resynchronization modules, this proposal has a built-in mechanism to automatically evaluate the attacking context and adjust the resynchronization accordingly. As a result, the robustness against valumetric attacks is advantageously preserved.

The proposed resynchronization approach currently defeats cropping attacks, possibly combined with rigid transforms and uniform scaling, or even shot noise that does not impact too many landmarks. When the cropping attack is combined with a valumetric attack such as noise addition, the performance of the decoder rapidly collapses. In future work, it may therefore be worth investigating alternate signature definitions to cope with such cases and thereby address real-life scenarios such as, e.g., print-and-scan attacks of 3D objects. Another avenue for future investigations relates to the security of the proposed resynchronization. The inserted landmarks are expected to be hard to identify without access to the secret key that defines the quantization lattice in the signature space. The task of the adversary is indeed complicated by the fact that watermarked signatures cannot be observed independently. As a result, conventional security attacks against quantization schemes are difficult to apply [18]. However, it is still unclear if the prior knowledge about the secret landmarks being placed onto a known reference shape can provide an edge to the adversary.

6. REFERENCES

- [1] Kai Wang, Guillaume Lavoué, Florence Denis, and Atilla Baskurt, "A comprehensive survey on three-dimensional mesh watermarking," *IEEE Transactions on Multimedia*, vol. 10, no. 8, pp. 1513–1527, December 2008.
- [2] Oliver van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or, "A survey on shape correspondence," *Computer Graphics Forum*, vol. 30, no. 6, pp. 1681–1707, September 2011.
- [3] Kai Wang, Guillaume Lavoué, Florence Denis, and Atilla Baskurt, "Three-dimensional meshes watermarking: Review and attack-centric investigation," in *Proceedings of the International Workshop on Information Hiding*, June 2007, vol. 4567 of *Lecture Notes in Computer Sciences*, pp. 50–64.
- [4] Ingemar J. Cox, Matthew L. Miller, Jeffrey A. Bloom, Jessica Fridrich, and Ton Kalker, *Digital Watermarking and Steganography*, Morgan Kaufmann Publishers Inc., 2nd edition, 2007.
- [5] Y.-P. Wang and S.-M. Hu, "A new watermarking method for 3D models based on integral invariants," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, no. 2, pp. 285–294, March–April 2009.
- [6] Patrice Rondao Alface and Benoit Macq, "Blind watermarking of 3D meshes using robust feature points detection," in *Proceedings of the IEEE International Conference on Image Processing*, September 2005, vol. I, pp. 693–696.
- [7] Patrice Rondao Alface, Benoit Macq, and François Cayre, "Blind and robust watermarking of 3D models: How to withstand the cropping attack?," in *Proceedings of the IEEE International Conference on Image Processing*, September 2007, vol. V, pp. 465–468.
- [8] Roland Hu, Li Xie, Huimin Yu, and Baocang Ding, "Applying 3D polygonal mesh watermarking for transmission security protection through sensor networks," *Mathematical Problems in Engineering*, January 2014.
- [9] Jae-Won Cho, Rémy Prost, and Ho-Youl Jung, "An oblivious watermarking for 3-D polygonal meshes using distribution of vertex norms," *IEEE Transactions on Signal Processing*, vol. 55, no. 1, pp. 142–155, January 2007.
- [10] Kai Wang, Guillaume Lavoué, Florence Denis, and Atilla Baskurt, "Robust and blind mesh watermarking based on volume moments," *Computer Graphics*, vol. 35, no. 1, pp. 1–19, February 2011.
- [11] Xavier Rolland-Nevière, Gwenaél Doërr, and Pierre Alliez, "Triangle surface mesh watermarking based on a constrained optimization framework," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 9, pp. 1491–1501, September 2014.
- [12] Ivan Sipiran and Benjamin Bustos, "Harris 3D: A robust extension of the Harris operator for interest point detection on 3D meshes," *The Visual Computer*, vol. 27, no. 11, pp. 963–976, November 2011.
- [13] The MathWorks Inc., "Optimization toolbox user's guide," 2013, <http://www.mathworks.com/access/helpdesk/help/toolbox/optim/>.
- [14] Massimiliano Corsini, Mohamed-Chaker Larabi, Guillaume Lavoué, Oldrich Petrík, Libor Vása, and Kai Wang, "Perceptual metrics for static and dynamic triangle meshes," *Computer Graphics Forum*, vol. 32, no. 1, pp. 101–125, February 2013.
- [15] Guillaume Lavoué, Elisa Drelie Gelasca, Florent Dupont, Atilla Baskurt, and Touradj Ebrahimi, "Perceptually-driven 3D distance metrics with application to watermarking," in *Applications of Digital Image Processing XXIX*, August 2006, vol. 6312 of *Proceedings of SPIE*.
- [16] Martin A. Fischler and Robert C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, June 1981.
- [17] Brian Chen and Gregory W. Wornell, "Quantization index modulation: A class of provably good methods for digital watermarking and information embedding," *IEEE Transactions on Information Theory*, vol. 47, no. 4, pp. 1423–1443, May 2001.
- [18] Luis Pérez-Freire, Pedro Comesaña, and Fernando Pérez-Gonzalez, "Information-theoretic analysis of security in side-informed data hiding," in *Proceedings of the International Workshop on Information Hiding*, June 2005, vol. 3727 of *Lecture Notes in Computer Science*, pp. 131–145.