# SPREAD TRANSFORM AND ROUGHNESS-BASED SHAPING TO IMPROVE 3D WATERMARKING BASED ON QUADRATIC PROGRAMMING

*Xavier Rolland-Nevière, Gwenaël Doërr*

Technicolor R&D France

*Pierre Alliez*

Inria Sophia Antipolis - Mediterranée

## ABSTRACT

Modulating the distances between the vertices and the center of mass of a triangular mesh is a popular approach to watermark 3D objects. Prior work has formulated this approach as a quadratic programming problem which minimizes the geometric distortion while embedding the watermark payload in the histogram of distances. To enhance this framework, we introduce two watermarking components, namely the spread transform and perceptual shaping based on roughness information. Benchmarking results showcase the benefits of these add-ons with respect to the fidelity-robustness trade-off.

***Index Terms***— 3D watermarking, spread transform, roughness

## 1. INTRODUCTION

Watermarking consists in imperceptibly altering multimedia content to transmit a message. This technology is instrumental in content protection systems due to the ability of the embedded watermark to survive subsequent modifications [1]. 3D models have become ubiquitous in many industrial applications such as movie production, video games and computational engineering. The need for watermarking methods dedicated to the protection of 3D assets has thus increased.

Three challenging issues must be addressed in 3D watermarking. First, meshes are irregular samplings of surfaces. As a result, signal processing tools, which are routinely used for uniformly-sampled signals such as audio and/or video cannot be readily used. Second, various defects, e.g. self-intersections, are commonly found in 3D models and practical algorithms will need to deal with such imperfect inputs. Third, meshes can undergo a large variety of alterations which may be hard to tackle. For these reasons, designing blind and robust 3D watermarking methods is still a scientific challenge. Indeed, without the original content at detection, watermarking methods cannot rely on recent advances in shape matching to perform registration. The synchronization between the embedder and decoder, and the robustness of the watermark is therefore usually achieved by using very resilient 3D primitives.

A popular 3D watermarking technique consists in modulating the distances between the vertices and the center of mass of a triangular mesh. In Section 2, we first review the different evolutions around this watermarking paradigm prior to detailing a quadratic programming (QP) framework that has been used to formalize this approach in Section 3. We then introduce two add-ons to complement this baseline approach in Section 4: the spread transform and perceptual shaping based on roughness information. Benchmarking results reported in Section 5 clearly illustrate the added value of the introduced components with respect to the fidelity-robustness watermarking trade-off. Eventually, we discuss in Section 6 a number of avenues for follow-up investigations.

## 2. RELATED WORK

The embedding domain can be used to classify the variety of published 3D watermarking algorithms [2]. On the one hand, most of the transforms which have been considered so far for watermarking (spectral, multi-resolution) induce subsequent non-blind watermark decoding [3]. On the other hand, spatial-based approaches have been found to yield very stable embedding primitives for watermarking. In particular, a pioneer work modifies the distances between the mesh vertices and center of mass to convey information [4].

More specifically, the watermark carrier is the average of the radial distances inside each bin of the histogram of said distances. Information is transmitted by modulating this average value around some threshold. The robustness of this scheme is brought by: (i) the very large stability of the center of mass against many attacks such as noise addition; (ii) the use of radial distances, which are invariant to rigid transforms; and (iii) the computation of a scale invariant data-dependent histogram. This baseline algorithm has then been further improved with a number of variants. For instance, more integral features, e.g. area-weighted [5] or volume-weighted geometric measures [6], have been considered to reinforce robustness. Alternatively, fidelity can be improved by a post-process which mitigates the embedding distortion [7] or by locally adjusting the watermarking strength based on some information of roughness [8]. Many perceptual distortion metrics indeed rely on the fact that geometrical alterations are more perceptible in smooth regions than in rough ones.

However, the baseline system has two major shortcomings: the causality issue is not properly tackled, i.e. the watermarking operation can alter the center of mass, and the actual vertex relocation strategy relies on an ad-hoc 'histogram mapping function', which is essentially a power function whose parameter is iteratively determined so that the average radial distance in a bin reaches a target value. To tackle both issues, the watermarking process has been previously formulated as a QP problem [9] that is detailed in Section 3. In this paper, we plug in two well known watermarking components in this QP framework. We introduce (i) a spread transform (ST) in each bin in an attempt to diversify the solution space and (ii) roughness-dependent weights in the fidelity constraint to perceptually adjust the watermark power.

## 3. QUADRATIC PROGRAMMING FRAMEWORK

A triangular surface mesh $\mathcal{M}$ is defined by its set of $n_v$ vertices, and its set of faces and edges. The Cartesian coordinates of vertex $\mathbf{v}_i$, $(i \in [\![0, n_v - 1]\!])$ are $(x_i, y_i, z_i)$, and its spherical coordinates with respect to the center of mass $\mathbf{G}$ are denoted $(\rho_i, \theta_i, \phi_i)$. $\mathbf{u}_i$ denotes the unit radial direction vector (from $\mathbf{G}$ to $\mathbf{v}_i$). Let $\mathbf{m}$ denote the watermark payload with $n_b$ bits and $\alpha > 0$ the watermark embedding

strength. A superscript '$^w$' refers to a watermarked variable. Vectors are written in column layout by convention.

To construct the watermarking carrier, the set of radial distances $\boldsymbol{\rho} = \{\rho_i\}$ is dispatched in a histogram having $n_B$ bins. Its edges range from $\min(\boldsymbol{\rho})$ to $\max(\boldsymbol{\rho})$ and are evenly spaced by a step $s$. $\rho^j_{\min}$ denotes the lower boundary of bin $j$, and $N_j$ is the number of samples within it. For practical reasons, $\rho_i$ is normalized to $\tilde{\rho}_i \in [0,1]$ using the affine mapping $\tilde{\rho}_i = s^{-1}\left(\rho_i - \rho^{B_i}_{\min}\right)$, where $B_i$ denotes the index of the bin associated to the distance $\rho_i$. The payload $\mathbf{m}$ is then embedded by altering the average value inside the bins of the histogram. The extremal bins are never modified (i.e., $n_B = n_b + 2$) to preserve the upper and lower bounds of the histogram, but the resulting offsets are omitted in the remainder of the paper (i.e. $n_B = n_b$ afterwards). To embed a bit $m_j$, the average value of the normalized radial distances inside bin $j$ is raised above $0.5 + \alpha$ if $m_j = 1$, and lowered below $0.5 - \alpha$ otherwise.

The embedding process can then be formulated as a QP problem, where the objective is to find the arrangement of vertex radial displacements which minimizes some cost function subject to the constraints of (i) embedding the desired payload in the histogram of radial distances and (ii) addressing the causality issue [9]. State-of-the-art solvers can be used to find such optimal solutions and the watermarked mesh is finally obtained by adding the relocation vectors $s\Delta\tilde{\rho}^w_i \mathbf{u}_i$ to the vertex positions, where $\Delta\tilde{\rho}^w_i = \tilde{\rho}^w_i - \tilde{\rho}_i$ are the displacements given by the optimal QP solution.

To minimize the geometrical distortion, the cost function in the QP framework is set to the squared error metric (SE), given by $\sum_{i=0}^{n_v-1} \|\Delta\tilde{\rho}_i\|^2$. Let $\mathbf{W} \in \mathbb{R}^{n_B \times n_v}$ denote the matrix whose coefficients are $W_{j,i} = {N_j}^{-1}\delta_{j,B_i}$, where $\delta$ is the usual Kronecker delta, $\mathbf{t} \in \mathbb{R}^{n_B}$ a constant vector with entries 0.5 and $\mathbf{M} \in \mathbb{R}^{n_b \times n_b}$ a diagonal matrix containing the (antipodal) payload bits. The watermarking constraints are expressed as a set of linear inequalities:

$$\mathbf{M}\left(\mathbf{t} - \mathbf{W}\tilde{\boldsymbol{\rho}}\right) + \alpha < \mathbf{M}\mathbf{W}\boldsymbol{\Delta}\tilde{\boldsymbol{\rho}}^w. \tag{1}$$

To perfectly reconstruct the histogram of radial distances at decoding, changes are prohibited in the position of the center of mass ($\mathbf{G}^w = \mathbf{G}$), and the mapping between vertices and bins ($\forall i \in [\![1, n_v]\!]$, $B^w_i = B_i$). $\mathbf{G}$ is computed as the average vertex position, which is a linear combination of the Cartesian coordinates. The stability of $\mathbf{G}$ can thus be written with the derivatives of the spherical coordinates:

$$\sum_{i=0}^{n_v-1} \Delta\tilde{\rho}^w_i \begin{pmatrix} \cos\theta_i \cos\phi_i \\ \sin\theta_i \cos\phi_i \\ \sin\phi_i \end{pmatrix} = \mathbf{0}. \tag{2}$$

Geometrically, it means that the sum of all vertex displacements is null. Finally, the mapping between bins and vertices is guaranteed by setting bounds for the unknowns: $-\tilde{\rho}_i < \Delta\tilde{\rho}^w_i < 1 - \tilde{\rho}_i$. Since the edges of the histogram are evenly spaced, and the extremal bins are not altered, these constraints are sufficient to formulate the watermarking problem.

At the receiver side, extracting the payload reduces to constructing the histogram of normalized radial distances $\tilde{\boldsymbol{\rho}}$ and comparing the average inside each bin to the threshold value $0.5$.

## 4. CONTRIBUTIONS

### 4.1. Spread-Transform

ST is routinely used in watermarking and amounts to projecting the carrier signal onto a pseudo-random sequence $\mathbf{p}$, i.e. to computing the inner product between the two sequences, prior to applying some embedding algorithm. The spreading sequence typically has zero-mean and unit norm with samples drawn from either a normal or uniform distribution. If the watermark modulation is binary, it is simply a regular spread-spectrum watermarking [10]. If the embedding mechanism instead relies on binning, it is the so-called spread-transform dither modulation (STDM [11]). While the latter has great potential theoretically, its practical performances are closely related to the distribution of the carrier signal.

Prior work in 3D watermarking considered applying STDM directly to the radial distances of the vertices [8]. In this case, the watermarking system is highly sensitive to the slightest changes in the ordering of the distances in $\boldsymbol{\rho}$. To avoid such instability, one may be tempted to apply STDM to the average radial distances considered in the baseline QP framework to benefit from the stability inherited from the integration within the bin. However, empirical observations revealed that the distribution of these values is highly concentrated around 0.5, thereby making them ill-suited for binning schemes. For large quantization steps, a single bin is used; for small quantization steps, robustness performances are worse than for the baseline system. Still, it may be useful to keep the ST in an attempt to diversify the solution space of the QP framework. The system is then quite close to spread spectrum except that the individual displacement for each vertex is not given by a generic equation but is instead driven by the solver that optimizes the QP problem.

Adding ST to the baseline QP framework modifies the watermark constraints in the embedding process. First, the number of bins in the histogram is multiplied by the spreading factor $k$, i.e., $n_B = k.n_b$ (omitting the extremal bins at both ends), and the average normalized radial distances are partitioned into $n_b$ consecutive carrier sequences with $k$ values. Let $\tau = 0.5\sum_{i=0}^{k-1} p_i$ denote the projection of $\mathbf{t}$ onto the spreading sequence. Projecting each carrier sequence onto $\mathbf{p}$ yields $n_b$ values $c_i$ and, depending on the bit $m_j$, $c_j$ is either raised above $\tau + \alpha$ or lowered below $\tau - \alpha$. This translates into a revised QP watermark constraint:

$$\mathbf{M}\boldsymbol{\Psi}\left(\mathbf{t} - \mathbf{W}\tilde{\boldsymbol{\rho}}\right) + \alpha < \mathbf{M}\boldsymbol{\Psi}\mathbf{W}\boldsymbol{\Delta}\tilde{\boldsymbol{\rho}}^w, \tag{3}$$

where $\boldsymbol{\Psi}$ denotes a $n_b \times n_B$ matrix, whose rows contain $k$ values of the spreading sequence $\mathbf{p}$, shifted by a multiple of $k$. For $k = 1$, $\boldsymbol{\Psi}$ is the identity matrix, and the constraint simplifies to Equation (1).

At decoding, the averages inside the $n_B$ bins of the histogram are computed, the $n_b$ sequences are projected back onto $\mathbf{p}$ and the resulting values are compared to $\tau$.

### 4.2. Roughness-driven Perceptual Shaping

Solutions of the QP problem tend to relocate all the vertices in a bin either outwardly or inwardly (depending on the bit value) with respect to the center of mass. This creates noticeable ring-like distortions on the mesh surface. Such artifacts are not adequately captured by the squared error metric, but perceptually-driven metrics such as the MSDM [12] are substantially more sensitive to them [13]. This calls for replacing the SE metric in the QP framework with another one that would be more aligned with human perception. Unfortunately, most of the existing perceptual metrics cannot be written as a quadratic function in $\boldsymbol{\Delta}\tilde{\boldsymbol{\rho}}$ (without large approximations) and, as such, cannot be easily integrated within the QP framework [14].

This being said, it is straightforward to scale the individual terms of the SE cost function by some weights $\mathbf{w} = \{w_i\}$ in an attempt to obtain a perceptually-driven weighted error $\text{wSE} = \sum_{i=0}^{n_v-1} w_i \|\Delta\tilde{\rho}_i\|^2$. In particular, based on previous findings, it makes sense to tie these weights to the local roughness in order to harden the fidelity constraint in smooth areas of the object and,

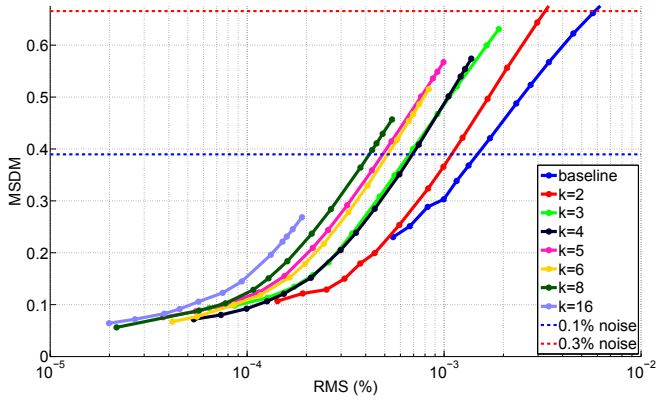


**Fig. 1**. Median RMS vs. median MSDM: assessing the embedding distortion when the embedding strength $\alpha$ is varied.

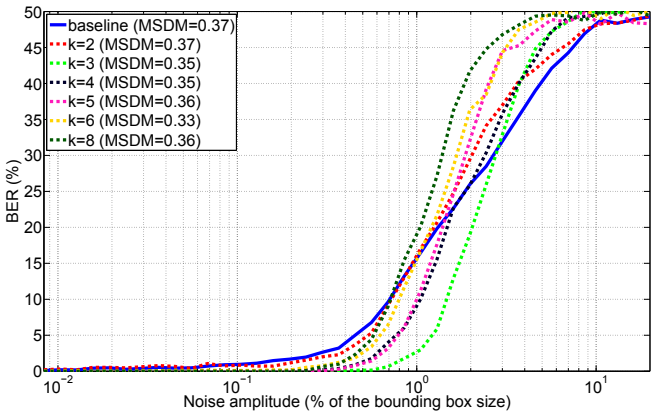


**Fig. 2**. Average BER against noise attacks (amplitude as a ratio of the size of the bounding box) for different spreading lengths and a target MSDM of 0.37.

conversely, to relax it in rough regions [8]. In this paper, the local roughness $r_i$ is estimated at each vertex and is derived from statistics relating to the principal curvatures [15]. Empirically, these roughness values have a distribution with few large outliers and small standard deviation and therefore need to be post-processed. First, outlying values are clipped to a minimal and maximal threshold, set to $5\%$ of the lowest and largest values. Next, all values are (affine) mapped to obtain $\tilde{\mathbf{r}} = \{\tilde{r}_i\}$ in $[0, 1]$. Finally, the weights $\mathbf{w}$ are set to $\mathbf{1} - \tilde{\mathbf{r}}$.

To ensure that vertices are not relocated to arbitrarily large distances, the QP cost function is defined as a linear trade-off between the SE and the wSE metric. Denoting by $\lambda \in [0, 1]$ the control parameter, the cost function is simply given by $(\mathbf{1} - \lambda\tilde{\mathbf{r}}) \cdot \boldsymbol{\Delta}\tilde{\boldsymbol{\rho}}^2$. For $\lambda = 0$, the cost function reduces to the SE metric (baseline system); for $\lambda = 1$, the fidelity constraint corresponds to the wSE metric.

## 5. EXPERIMENTAL RESULTS

For benchmarking purposes, we considered a database of 10 meshes having between 20k and 100k vertices [16]. For each 3D model, seven watermark payloads of $n_b = 16$ bits are randomly generated and embedded, thereby creating creating 70 differently watermarked meshes. Multiple spreading lengths are surveyed, including $k = 1$ (baseline), 2, 3, 4, 5, 6, 8, and 16. Note that using a fixed payload size means that the number of bins in the histogram of radial distances increases with $k$.

### 5.1. Embedding Distortion with ST

The embedding distortion is recorded for multiple embedding strengths $\alpha \in [0.001, 0.4]$, knowing that large $\alpha$ values are likely to create unsolvable QP problems in the baseline framework. Two distortion metrics are included in the benchmark: the root mean square error (RMS) and the MSDM, which is expected to be more in accordance to human perception. The first metric pools the Euclidean distances between vertices in $\mathbb{R}^3$ and is expressed as a ratio of the size of the bounding box for normalization purposes. The second metric is asymmetric and we report the distance from the original to the watermarked mesh. The median results are reported in Figure 1.

Increasing the spreading length decreases the lower bound of achievable embedding distortions for both metrics. However, if a target RMS can be reached with a spreading length $k$, increasing this value yields larger perceptual distortion, as measured by MSDM. This is due to the fact that increasing the spreading length translates into a larger number of bins in the histogram. As a result, the relocation energy is spread on more, yet slightly smaller, ring-like perturbations, which further triggers the MSDM. Conversely, for a target MSDM value, increasing $k$ lowers the RMS. Indeed, to maintain the MSDM while increasing the number of ring artifacts, it is necessary to greatly reduce their amplitudes and thus to decrease the RMS. All in all, these observations corroborate that the MSDM is more sensitive to the ring-like embedding artifacts.

To place these first results into perspective, two dotted lines have been added in Figure 1 to indicate the MSDM distortion introduced by uniform noise addition for two different levels of noise. The lower level ($0.1\%$ noise amplitude with respect to the size of the bounding box) is barely noticeable, while the larger one ($0.3\%$) represents a perceptible alteration. This clearly highlights that ST is most useful to bring flexibility in the acceptable distortion range for the QP framework. Without ST, we would need to use very small embedding strengths to get in this region and the watermark would therefore have very little robustness.

### 5.2. Robustness with ST

By design, all methods are invariant to rigid transforms, reordering of vertices, and uniform scaling, and we therefore investigate the impact of ST on robustness after uniform noise addition. In order to guarantee a fair comparison, the embedding strength $\alpha$ has been adjusted for each spreading length $k$ in order to obtain a MSDM distortion close to 0.37. Figure 2 illustrates the recorded bit error rate (BER) when increasing the attacking strength. There is no curve for $k = 16$ since this setup consistently produces MSDM distortion lower than the target value and therefore exhibits extremely poor robustness.

On average, using a longer spreading sequence decreases the BER with respect to the baseline QP framework for noise levels lower than a $1\%$ cut-off threshold. Conversely, the opposite phenomenon appears for stronger attacks: the BER rockets up more quickly to $50\%$. In other words, ST preserves the watermark transmission quality longer but collapses more drastically when the attack exceeds the capabilities of the system. The same trend is observed at various target MSDM value. As a rule of thumb, the larger the target

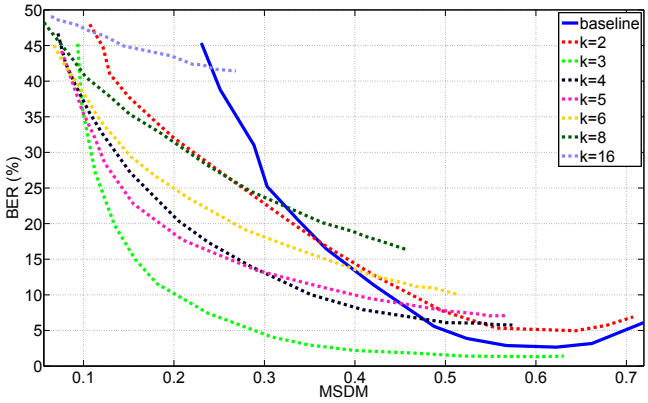


**Fig. 3**. Average BER vs. perceptual embedding distortion for $1\%$ uniform noise addition.

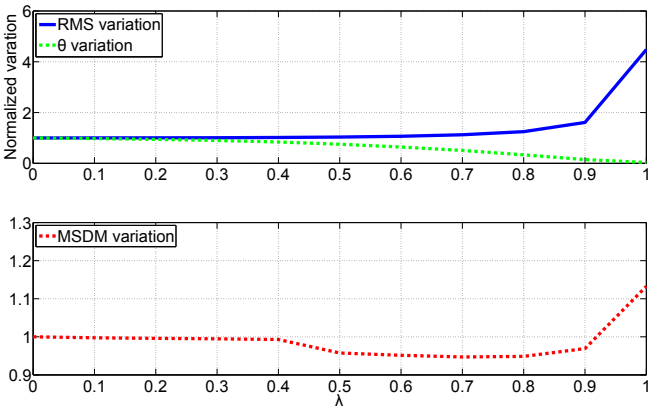


**Fig. 4**. Fidelity variation evaluation using three distortion metrics, depending on the parameter $\lambda$ which adjust the trade-off in the cost function between SE and the roughness-driven wSE.

MSDM is, the longer the system survives and the sharper the transition regime is. To provide a better understanding of the situation, Figure 3 depicts the robustness vs. fidelity trade-off for $1\%$ uniform noise addition.

The typical operating region for 3D watermarking systems is given by a MSDM in $[0.35, 0.4]$, to have the largest possible distortion that remains hardly noticeable. In this region, ST obviously manages to lower the BER for the same distortion level, and the spreading value $k = 3$ offers the best fidelity-robustness trade-off.

Due to space restriction, we do not report on other attacks such as quantization, smoothing, etc. Nevertheless, the same general trend is observed: for a given level of fidelity, ST provides a boost of robustness for lower levels of attack until a drop-off threshold where performances collapse faster than for the baseline QP 3D watermarking framework.

### 5.3. Embedding Distortion with Perceptual Shaping

To evaluate the impact of introducing perceptual weights, relating to roughness information, on fidelity, we measure the embedding distortion with three alternate metrics: the RMS, the MSDM, and an additional metric, defined as $\theta = \mathbf{w} \cdot \boldsymbol{\Delta\rho^2} \left( \|\boldsymbol{\Delta\rho^2}\| \|\mathbf{w}\| \right)^{-1}$. Geometrically, the last metric can be seen as the cosine of the angle between the vector of vertex relocation amplitudes and the vector of roughness-driven weights. Figure 4 depicts the variations of these metrics when varying the mixing parameter $\lambda$ from 0 (SE minimization) to 1 (only minimizing the wSE), for $\alpha = 0.05$ and without ST (baseline).

Essentially, the RMS is aligned with the SE cost function and $\theta$ is aligned with the wSE. As a result, the RMS increases with $\lambda$ since the underlying cost function in the QP optimization framework gradually moves away from SE. Conversely, raising $\lambda$ strengthens the perceptual constraints in the cost function and makes the relocation amplitude vector $\boldsymbol{\Delta\rho^2}$ deviate more and more from the weighting direction $\mathbf{w}$. Visually, this translates into the ring-like artifacts being attenuated in smooth regions and being amplified in the roughest parts of the 3D model. All these observations are matching the expected behavior of the system and the interesting plot is actually the one reporting the variation of the MSDM metric. Due to its complex definition, this metric cannot be easily integrated within the QP framework. Still, the MSDM considers roughness information among other distortion cues to evaluate fidelity. As a result, introducing a weighting scheme accounting for this information in the QP framework appears to provide a means to decrease the MSDM distortion. Indeed, there seems to be an optimal mixing parameters $\lambda^*$ for which the MSDM is $7\%$ lower than with the baseline QP framework ($\lambda = 0$). A quick sanity check confirms that this fidelity gain does not induce any noticeable change in robustness. In other words, introducing this roughness-driven perceptual shaping strategy in the QP framework further improves the fidelity-robustness trade-off.

## 6. CONCLUSION

In this paper, we introduced two add-ons to enrich the QP framework for watermarking the radial distances of a surface mesh. The ST leverages the advantages of the stability provided by the integration within each bin of the histogram and the diversity of the solution space provided by the subdivision of the bins. This variant is shown to improve the quality of watermarking transmission: the robustness is improved up to a cut-off attacking strength where robustness performances collapse faster than with the baseline system. The integration of roughness-driven perceptual weights in the QP framework also provides a means to account for complex fidelity metrics such as the MSDM, that cannot be directly incorporated into the framework without compromising robustness. Although not reported in this paper, these two add-ons can be successfully combined to further improve the fidelity-robustness watermarking trade-off.

In future work, we will investigate how to gain greater control over the embedding distortion by using more sophisticated cost functions that cannot be reduced to simple wMSE, e.g. cost functions having cross terms in their formulation to account for the interdependency between neighbor vertices. Moreover, based on past experience with audio-visual content, we will explore the possibility to model the relocation directions themselves from perceptual principles rather than simply scaling the relocation amplitude. Finally, we notice that a by-product of the ST component is that it introduces a secret in the watermarking framework. We will evaluate its impact on 3D watermarking security.

## 7. REFERENCES


[1] 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.

[2] 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.

[3] Ming Luo, *Robust and Blind 3D Watermarking*, Ph.D. thesis, University of York, UK, May 2006.

[4] 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.

[5] Patrice-Rondao Alface, Benoit Macq, and Franois Cayre, "Blind and robust watermarking of 3D models: How to withstand the cropping attack?," in *Proceedings of the IEEE International Conference on Image Processing*, October 2007, vol. V, pp. 465–468.

[6] 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.

[7] Adrian G. Bors and Ming Luo, "Optimized 3D watermarking for minimal surface distortion," *IEEE Transactions on Image Processing*, vol. 22, no. 5, pp. 1822–1835, May 2013.

[8] Rony Darazi, Roland Hu, and Benoit Macq, "Applying spread transform dither modulation for 3D-mesh watermarking by using perceptual models," in *Proceedings of the IEEE International Conference on Acoustics Speech and Signal Processing*, March 2010, pp. 1742–1745.

[9] Roland Hu, Patrice Rondao-Alface, and Benoit Macq, "Constrained optimisation of 3D polygonal mesh watermarking by quadratic programming," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, April 2009, pp. 1501–1504.

[10] Ingemar J. Cox, Joe Kilian, Thomson Leighton, and Talal Shamoon, "Secure spread spectrum watermarking for multimedia," *IEEE Transactions on Image Processing*, vol. 6, no. 12, pp. 1673–1687, December 1997.

[11] Brian Chen and Gregory W. Wornell, "Quantization index modulation: A class of provably good methods for digital watermarking and information embedding," *IEEE Transaction on Information Theory*, vol. 47, no. 4, pp. 1423–1443, May 1999.

[12] 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*, p. 63120L.

[13] Guillaume Lavoué and Massimiliano Corsini, "A comparison of perceptually-based metrics for objective evaluation of geometry processing," *IEEE Transactions on Multimedia*, vol. 12, no. 7, pp. 636–649, November 2010.

[14] The MathWorks Inc., "Optimization toolbox user's guide," 2013, `http://www.mathworks.com/access/helpdesk/help/toolbox/optim/`.

[15] Guillaume Lavoué, "A roughness measure for 3D mesh visual masking," in *Proceedings of the Symposium on Applied Perception in Graphics and Visualization*, July 2007, pp. 57–60.

[16] Kai Wang, Guillaume Lavoué, Florence Denis, Atilla Baskurt, and Xiyan He, "A benchmark for 3D mesh watermarking," in *Proceedings of the IEEE International Conference on Shape Modeling and Applications*, June 2010, pp. 231–235.