Digital Watermarking for Volumetric Datasets

Leonid I. Dimitrov and Miloš Šrámek

Commission for Scientific Visualization,
Austrian Academie of Sciences,
Vienna, Austria

Abstract

An idea for digitaly watermarking volumetric datasets (VDS) and preliminary results from a test implementation are presented. It consists of a combination of filtering and rendering operations on volumetric data allowing for identification of the voxels suitable for hiding the watermark and spreading the information in the gradients of the selected voxels. Thus, a secure and robust method for watermarking VDS could be developed which takes into account the specific properties of volumetric data and advantageously utilizes them.

1. Introduction

Today's medical imaging modalities—CAT, MRI, fMRI, USG, PET, SPECT, EEG, MEG, etc., deliver with ever-increasing speed a multitude of volumetric datasets each covering a different aspect of the complex nature of human anatomy and physiology. One of the problems associated with the acquisition and storage of such datasets concerns the obtainment and attainment of the privacy of the acquired data. While all other medical records are strictly required to be kept private and protected against unauthorized manipulations, by far not the same level of scrutiny is applied to volumetric data, although they are of even greater sensibility and private nature. There is a stark contrast here between legal requirements and technical possibilities.

Still another area where authenticity and protection of volumetric data has to be achieved is found in the emerging field of volume graphics. Here, methods and algorithms are being developed aiming to augment and perhaps even replace one day the older 3D modeling primitives like polygons and solids by a newer and more promising approach: model and represent the objects as volumetric datasets. Such an approach would free the modeling and rendering systems from the necessity to deal with different object representations, allow for a simplified and unified such representation and increase the efficiency. Besides, a plentitude of novel modeling techniques could be applied to such data and provide for more diversity and exciting new possibilities. Clearly, technical means for protecting the VDS author's and owner's rights and enforcing the authenticity and uniqueness of the created datasets have to be developed and applied here, too.

A similar situation existed not long ago in the area of computer acquired and/or generated 2D imagery or 3D models. Meanwhile, appropriate computer-based methods and algorithms were developed targeting the unique, robust, undestroyable and imperceptible marking of those objects in such a way as to ensure their authenticity and protect them against unauthorized use. These methods are similar to the older and well-known watermarks used to protect bank notes and other valuable papers hence their collective description as *digital watermarking*[1, 2, 3, 5].

All digital watermarking (DW) methods function in such a way that they exploit the inherent redundancy in the data and try to hide additional information identifying the rightful owners and inevitably hereby altering the original objects (images or 3D meshes). Good watermarking methods do this without losing too much valuable original information or making the alterations too obvious and easy to remove. This is mostly done by hiding the information identifying the object and/or its author in the unavoidable noise component or least significant areas of the original object/image [12]. A couple of conditions need to be met in order to make this approach work: the changes have to be unobtrusive and insignificant to the original information content (imperceptible), extractable only by the authorized users (secure), and hard to remove (robust) [6, 13].

Volumetric imaging has reached today a state where digital means for enforcing the authenticity and privacy of the generated data are urgently needed, too. Unfortunately, a specific digital watermarking method for volumetric datasets has not yet been developed to the best of our knowledge. Clearly, the well-known methods from the 2D imaging area could be applied here, too, with all their weaknesses and drawbacks, but they weren't specifically developed for this kind of application and consequently can't reach the degree of suitability required. On the other hand, volumetric data offer some unique features which could be advantageously used for developing efficient, robust and secure digital watermarking methods specifically targeted at this kind of data and exploiting their unique properties. This is the motivation for conducting this research.

2. Methodology and Design

The selected methodology for solving our problem is easily pointed out: we want to draw upon the collective knowledge in the field of information hiding and develop methods which take into account the meanwhile established general principles in that scientific area and develop new methods for digital watermarking which target specifically volumetric datasets and build upon the unique features they exhibit.

Digital watermarking exists as a sub-discipline in the more general information hiding field. All methods for digital watermarking hide additional information, viz the watermark, in the data to protect. This has to be accomplished in a *secure* and *robust* way. Secure means in the first place that only authorized recipients—men or machines—should be able to extract the additional information. Methods for achieving exactly this are studied by cryptography and mean in general that a strong, key-protected symmetric or asymmetric cryptographic method has to be applied. We can safely assume that such methods exist and can be applied to our problem. But *secure* in the context of steganography to which digital watermarking belongs means further that even the mere fact of the presence of additional information has to be undetectable. This is a very strong requirement—the authorized recipient has to be able to detect the presence of the watermark finally, too—and is replaced by the weaker requirement that the additional information has to be *imperceptible* i.e. to human beings and undetectable by "common processing" for passive men-in-the-middle attackers.

The human vision (or auditory) system is easy to deceive (cf. e.g. [26]) but to hide information (e.g. digital watermarks) in other information sources without noticeably altering their statistical properties is virtually impossible or at least very difficult. The situation with digital watermarking is even worse since the information addition has to happen also *robustly*, i.e. it should be impossible or, at least very hard for unauthorized adversaries (active men-in-the-middle attackers) to destroy or even remove the watermark. Active adversaries are normally assumed to dispose of unlimited resources—time and computing power—and none of the presently known digital watermarking methods is strong enough to survive the efforts of a dedicated adversary, but again we have to revise our requirements and want to assume that the only operations which a digitally watermarked volumetric dataset has to be able to safely survive should be limited to common filtering, cropping and linear transformation operations (cf. [20]). This looks at first like a strong limitation, but considering that everything else would alter the volumetric data beyond usability, too, shows that it is as such outside of the scope of our efforts.

Bearing in mind the above considerations, we could tackle our problem firstly by simply trying out the well-established digital watermarking methods, e.g. bit substitution techniques [7]. The simplest way would be to "hide" the watermark by substituting the least significant bits (LSBs) of the consecutive voxels by the next bits of the watermark. Real world volumetric data (radiological sources like CAT or MRI) are subject to considerable pollution by noise which usually shows as more or less random minute fluctuations in the LSBs. This means in general that we could conduct the above substitution without losing any relevant information and rendering the volumetric data useless. On the other hand though, such a straight-forward approach is by no means secure—everybody could detect the presence of the watermark by simply looking at the distribution/visual pattern of the LSBs which will differ significantly from the rest of/other VDS of the same source. Robustness is also not a prominent feature of this over-simplistic method—the watermark would survive cropping and even some linear operations like scaling to some extent, but even very "mild" filtering operations like smoothing with a Gaussian kernel will inevitably destroy it.

The above example wasn't really meant as a serious attempt at digital watermarking, but it is a good starting point for further considerations because it shows us already what the weak points of a digital watermarking scheme could be and how we could try to circumvent them. Firstly, we have to make a future DW method more secure, e.g. by distributing the DW bits pseudorandomly over the whole VDS. Doing this without taking care of the cases where already used positions are addressed again imposes the danger of overwriting previously hidden bits (collisions). Provided, we generate the hidden bit positions by employing a cryptographically secure pseudorandom number generator (PRNG), we could seed it with a secret number and use that as key which has to be transmitted to the recipient by a secure key-exchange protocol, of course. In most cases though, the recipient is the rightful owner or author of the VDS in question and only he has to be able to check the presence or absence of the DW, so no real transfer and hence no immediate danger of compromising the key exists.

Pseudorandomly distributing the DW bits over the whole VDS is just the first step. Next, the sequence of hidden bit positions has to be restorable which means that either no collisions happen, or that they are tolerable, or correctable, e.g.

through the use of an error-correcting code. In order not to destroy the original statistical properties of the VDS used as cover, we can either use DWs whose size is negligible compared to that of the VDS, or try and model the original VDS bit distribution in the distribution of the DW bits to hide.

Through a combination of the above tools, we could achieve any degree of security but, alas, at the cost of robustness. We make it hard for a passive men-in-the-middle attacker to detect the hidden DW, but at the same time it becomes more vulnerable to malicious active attackers because even without a-priori knowledge about the presence of a DW, it would be possible for them to destroy or even entirely remove it through a sequence of permitted operations: filtering, cropping, scaling. In order to attain robustness, we need more, e.g. distribute the DW bits in higher than the LSBs of the VDS or in other than the spatial domains. It remains to be researched in how far this can be done without destroying the valuable VDS information and how to identify areas and voxels suitable as carriers of the DW bits. Such an approach would perhaps increase the robustness against filtering attacks, but it is still not entirely clear how to undermine possible attacks by cropping or even linear transformations. At this stage, we have only certain ideas like building up the pseudorandom bit locations chain in a way independent of the VDS dimensions (increases the robustness against cropping attacks) or orientation (robustness against linear transformations).

One possible scenario could be as follows: suppose, we want to hide a bit of secret information pertaining to a position in the VDS (voxel). We generate somehow a direction which depends only on the contents of the voxel and not its location in the VDS, shoot a traversing ray in that direction, step a pseudorandom number of steps along the ray and hide the bit in the reached voxel. This approach would survive cropping and scaling attacks and could be made arbitrarily secure and robust, provided the PRNG and bit hiding operation employed are resp. secure and robust enough. Voxel traversal in connection with volume rendering has been studied for years, and we are confident that a DW method along these development lines could be devised.

Up till now, we left the special properties of VDS unattended to in our considerations. The ideas and attempts sketched above pertain to (almost) any digital source used as steganography cover. However, volumetric datasets exhibit special properties, e.g. their sheer size or adequacy for visual interpretation, which could be advantageous to the construction of secure and robust specialized digital watermarking methods.

One idea concerns the use of gradient filters for encoding/hiding watermarks. The secret information (DW) can be hidden in systematic gradient changes during the DW embedding phase, i.e. minute changes are applied to the gradients in certain locations, identified through one of the above sketched cryptographically secure methods. This scheme amounts to a special gradient splatting operation. The gradient changes/perturbations cause later clearly visible/extractable pre-defined distortions (patterns) in images reconstructed/visualized in the DW extraction phase according to a specially tailored rendering and gradient estimation method. Constructing such a specialized gradient splatting operation means hiding the information in the frequency domain and hence achieves robustness without (much) additional efforts [8]. Such a scheme would exploit the visual characteristics of VDS and would leave us with a VDS targeting digital watermarking method.

Visually extracting/visualizing the DW encoded in the gradients is just one possibility. We could see this as a first stage in a two stage DW extraction method. The DW embedded in the encoding phase could be any information worth hiding including encrypted and error-code encoded bitstreams. In a second extraction phase it could be extracted as usual, decoded, decrypted and used in whatever way intended, e.g. as copy protection mechanism.

3. Results

We conducted a preliminary research in the above sketched directions with following (very preliminary) results:

The design decision we took was to develop a DW system for volumetric datasets which would hide the information (watermark) in a derived feature of the VDS in question, viz the voxel gradients which are usually used as a normal substitute in a subsequent visualization/rendering step. The embedding of the watermark should happen *imperceptibly, securely and robustly* in accordance with the requirements stated in e.g. [1]. Choosing the gradient as carrier furthers the robustness, since it spreads the information in the frequency transform domain. The watermark itself could be of any kind (any bitstream) but in a first approach we decided to use 3D patterns, themselves expressed/voxelized as VDS, since that would enable a direct and immediate visual extraction check of the watermark. By taking this decision, we don't give up on any generality or application, since the extraction process we implement is just one of many possible. Its important phase is the first one: reliable localization and extraction of the encoded gradient positions. The second, decoding phase could be anything in fact: at the development stage we opt for visualization, but it could be also something else, e.g. bit stream extraction.

We demonstrate our results with following example:

Starting with the VDS of an MRI scan of the head, we can extract and visualize isosurfaces in that (cf. 1) using a gradient shading method as implemented by our VDS processing and rendering system VORTEX [29].

In a second step, we embed the VDS into another, artificially created (voxelized) VDS, consisting in our test case of a set of regularly positioned spheres. VORTEX provides

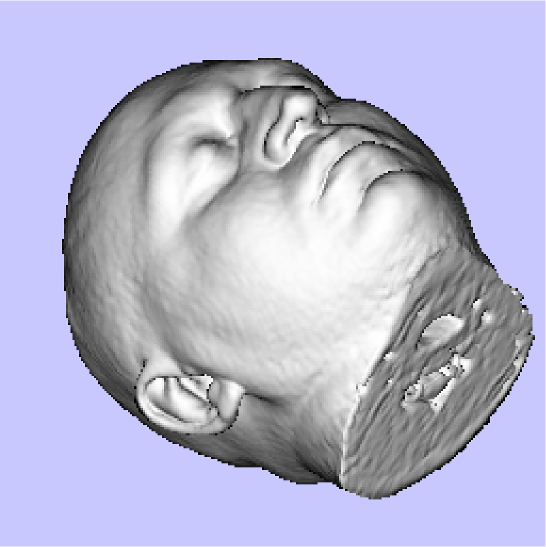

Figure 1: *Gradient-shaded isosurface: skin.*

the means for this operation, too, and Fig. 2 illustrates the resulting situation.

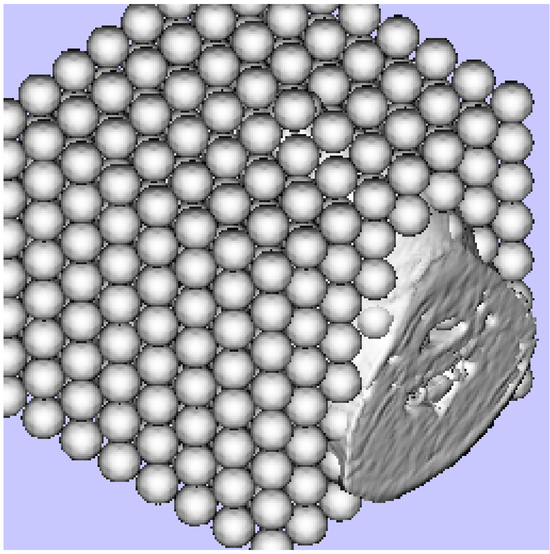

Figure 2: *VDS object embedded in an artificial VDS.*

We use this configuration then to identify all positions/voxels in the first VDS (head MRI) belonging to the desired isosurface (skin) AND intersected by the isosurface of the watermark VDS object (spheres). Next, we mark the gradients at all the intersection positions as "distorted" and hide this information (1 bit) at a pseudorandomly chosen secret position relative to the voxel in question. We seed the PRNG with a bit combination consisting of the isovalue, the voxel coordinates and a secret predefined initial seed. Besides, we discard all secret voxel positions lying on the isosurface itself and keep searching till we find a suitable one.

The resulting VDS exhibits no significant changes in grayvalue distribution under direct investigation, or isosurface visualization with "normal" gradient shading when used as usual, as Fig. 3 demonstrates.

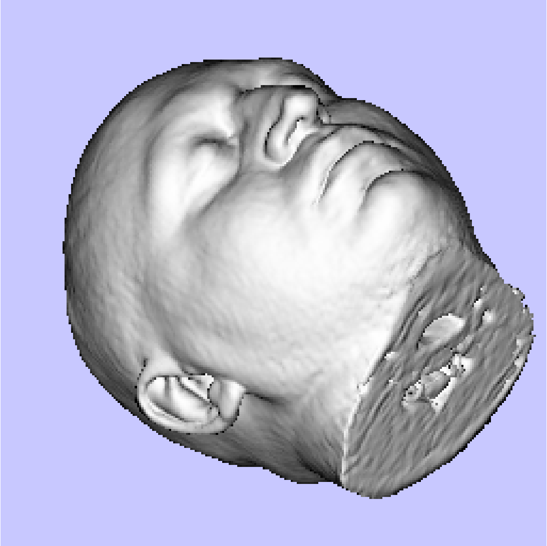

Figure 3: *Gradient-shaded isosurface in watermarked VDS.*

Also, the contrast enhanced difference image between the images in Fig. 1 and Fig. 3 exhibits no patterns or systematic changes as Fig. 4 shows.

Only when a "special" watermark/gradient-extraction scheme, seeded with the right seed/key, is applied, a clearly visible pattern, resulting from the embedding of the original VDS in the watermarking VDS, appears. Fig. 5 shows the results of trying this scheme a) on the original, "clean" VDS with some seed/key, b) on the watermarked VDS, but with a wrong seed/key, and c) on the watermarked VDS with the right key.

4. Conclusions

What we demonstrated is "quite" a secure VDS watermarking scheme—the embedded watermark is imperceptible and not easily detectable by statistical or other processing but by no means robust: even mild image processing operations would destroy the watermark entirely since the gradient distortion information was hidden in the LSBs of the securely chosen voxels. The same applies to cropping or affine operations: they destroy the positioning information since we address the chosen secret voxels em relatively to the current position, i.e. depending on the VDS dimension and orientation.

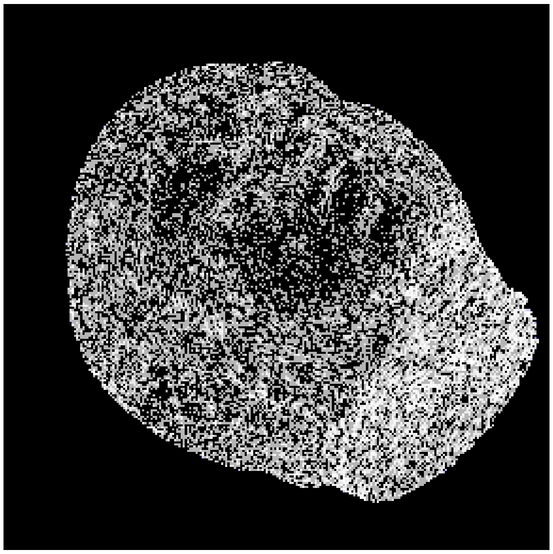

Figure 4: *Enhanced difference image between rendered isosurfaces.*

It should be clear though, that the demonstrated results are just a very first step in the direction we plan to go. They exhibit also enough potential for improvement and development.

Building upon the collective knowledge of the academic information hiding community and exploiting the special volumetric data characteristics, we intend to construct secure and robust digital watermarking methods for VDS which would hide the additional information/watermark in some spectral domain(s) in the form of systematic gradient changes. This information hiding should happen securely and robustly, cause no perceptible changes in the VDS in question and be reliably detectable by a corresponding DW extraction method. The method developed should enable the use of any information as digital watermark—from visible 3D objects or patterns to encrypted and encoded bitstreams.

References

1. Stefan Katzenbeisser and Fabien A. Petitcolas, editors. *Information hiding techniques for steganography and digital watermarking.* Artech House, MA, 2000. 1, 3

2. Steffen Moller, Andreas Pfitzmann, and Ingo Stierand. Rechnergestützte Steganographie: Wie sie Funktioniert und warum folglich jede Reglementierung von Verschlüsselung unsinnig ist. *Datenschutz und Datensicherheit*, 18(6):318–326, ???? 1994. 1

3. T. Aura. Practical invisibility in digital communications. In Ross Anderson, editor, *Information hiding: first international workshop, Cambridge, U.K., May 30–June 1, 1996: proceedings*, volume 1174 of *Lecture Notes in Computer Science*, pages 265–278, Berlin, Germany / Heidelberg, Germany / London, UK / etc., 1996. Springer-Verlag. 1

4. Jiri Fridrich. A new steganographic method for palette-based images. In *Proceedings of the Conference on Image Processing, Image Quality and Image Capture Systems (PICS-99)*, pages 285–289, Springfield, Virginia, April 25–28 1999. Society of Image Science and Technologie.

5. J. Zhao and E. Koch. Embedding robust labels into images for copyright protection. In *International Congress on Intellectual Property Rights for Specialised Information, Knowledge and New Technologies*, Vienna, Austria, 21–25 August 1995. 1

6. E. Koch and J. Zhao. Towards robust and hidden image copyright labeling. In I. (Ioannis) Pitas, editor, *Proceedings of 1995 IEEE Workshop on Nonlinear Signal and Image Processing (Neos Marmaras, Halkidiki, Greece, June 20–22, 1995)*, pages 452–455, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1995. IEEE Computer Society Press. 1

7. A. Z. Tirkel, G. A. Rankin, R. M. van Schyndel, W. J. Ho, N. R. A. Mee, and C. F. Osborne. Electronic watermark. In *Digital Image Computing, Technology and Applications (DICTA'93)*, pages 666–673, Macquarie University, Sidney, 1993. 2

8. R. Pickholtz, D. Schilling, and L. Milstein. Theory of spread spectrum communications - a tutorial, 1982. 3

9. L. M. Marvel, C. G. Boncelet, and C. T. Retter. Reliable blind information hiding for images. *Lecture Notes in Computer Science*, 1525:48–61, 1998.

10. Jian Zhao. Look, it's not there – digital watermarking is the best way to protect intellectual property from illicit copying. *BYTE*, January 1997.

11. Joshua R. Smith and Barrett O. Comiskey. Modulation and information hiding in images. In Ross J. Anderson, editor, *Information hiding: first international workshop*, volume 1174 of *Lecture Notes in Computer Science*, pages 207–226, Isaac Newton Institute, Cambridge, England, May 1996. Springer-Verlag, Berlin, Germany.

12. W. Bender, D. Gruhl, N. Morimoto, and A. Lu. Techniques for data hiding. *IBM Systems Journal*, 35(3&4):313–336, 1996. 1

13. Gerrit C. Langelaar, Jan C. A. van der Lubbe, and Reginald L. Lagendijk. Robust labeling methods for copy protection of images. In *Proceedings of SPIE Electronic Imaging '97, Storage and Retrieval for Image and Video Databases V*, San Jose, California, February 1997. 1

14. Martin Kutter, Frederic Jordan, and Frank Bossen. Digital signature of color images using amplitude modulation. In *Proc. SPIE Storage and Retrieval for Image and Video Databases*, volume 3022, pages 518–526, San Jose, California, 1997.

15. J. Ruanaidh, W. Dowling, and F. Boland. Phase watermarking of digital images, 1996.

16. K. Mantusi and K. Tanaka. Video steganography: How to secretly embed a signature in a picture, 1994.

17. J. Zhao. A www service to embed and prove digital copyright watermarks, 1996.

18. Neil F. Johnson and Sushil Jajodia. Exploring steganography: Seeing the unseen. *IEEE Computer*, 31(2):26–34, 1998.

19. Neil F. Johnson and Sushil Jajodia. Steganalysis of images created using current steganography software. In *Information Hiding*, pages 273–289, 1998.

20. J. Ruanaidh and T. Pun. Rotation, scale and translation invariant digital image watermarking, 1997. 2

21. H. Wang. An integrated progressive image coding and watermark system, 1998.

22. D. Kundur and D. Hatzinakos. A robust digital image watermarking method using wavelet-based fusion, 1997.

23. T. Wilson, S. Rogers, and L. Myers. Perceptual-based hyperspectral image fusion using multiresolution analysis, 1995.

24. Xian-Gen Xia, Charles G. Boncelet, and Gonzalo R. Arce. Wavelet transform based watermark for digital images. *Optics Express*, 3(12):497–511, 7 1998.

25. Tsekeridou and Pitas. Wavelet-based self-similar watermarking for still images.

26. Stephen E. Palmer. *Vision Science—photons to phenomenology*. The MIT Press, Cambridge, Massachusetts, USA, 1999. 2

27. I. Holländer, H. Petsche, L. I. Dimitrov, O. Filz, and E. Wenger. The reflection of cognitive tasks in EEG and MRI and a method of its visualization. *Brain Topography*, 9(3), 1997.

28. Leonid I. Dimitrov. Texturing 3D-reconstructions of the human brain with EEG-activity maps. *Human Brain Mapping*, 6(4), 1998.

29. L. I. Dimitrov and I. Holländer. A comparison of a classical and a novel volume rendering methods exemplified in a study of the human cortex. In H. U. Lemke, M. W. Vannier, K. Inamura, and A. G. Farman, editors, *Computer Assisted Radiology ’98. Excerpta Medica ICS1165*, page 888, Amsterdam, 1998. Elsevier. 3

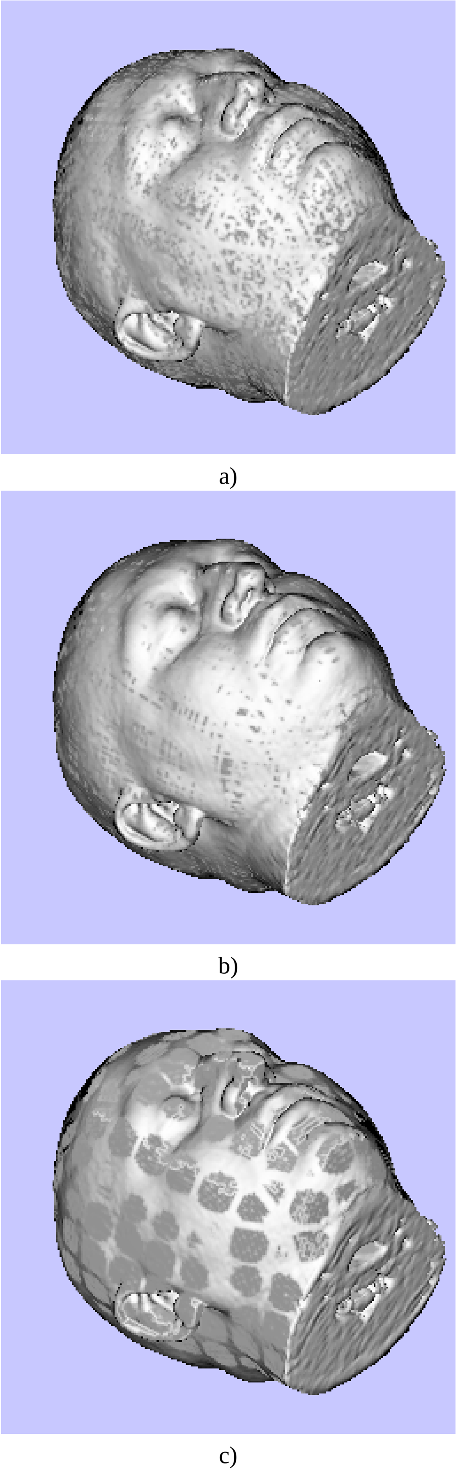

a)

b)

c)

Figure 5: *Results of watermark extraction/visualization.*