Authentication and Self-Correction in Sequential MRI Slices

Vassilis Fotopoulos · Maria L. Stavrinou ·
Athanasios N. Skodras

Published online: 14 October 2010
© Society for Imaging Informatics in Medicine 2010

Abstract One of the new challenges of Information Technology in the medical world is the protection and authentication of a variety of digital medical files, datasets, and images. In this work, the ability of magnetic resonance imaging (MRI) slice sequences to hide digital data is investigated and more specifically the case that the hidden data are the regions of interest (ROI) of the MRI slices. The regions of non-interest (RONI) are used as cover. The hiding capacity of the whole sequence is taken into account. Any ROI-targeted tampering attempt can be detected, and the original image can be self-restored (under certain conditions) by extracting the ROI from the RONI.

Keywords Medical imaging · MRI · ROI · Authentication · Self-correction · Integrity · JPEG2000

Background

Digital biosignals like digital radiological images and one-dimensional datasets are massively produced everyday in medical facilities around the world. The interoperability requirement for ubiquitous personalized health services has led to the creation of the electronic health record (EHR) and of the necessary exchange protocols. Generally, the EHR contains patient demographics, diagnostic reports or codes, prescriptions as well as radiological images. One of the most exciting future expectations for the EHR is its worldwide availability from any medical facility at any time.

In this context, healthcare systems have to upgrade from individually (personal or hospital) based towards more global e-health platform systems. These future personal e-health platforms, due to the strict ethics and legislative rules concerning medical data, have to ensure confidentiality, reliability (authenticity and integrity check), availability, and safe transfer of the medical files [1–3]. Apart from legislation issues, the quality of the biomedical images has to remain intact to ensure a correct diagnosis. Thus, an automated framework is needed on these e-health platforms to verify the authenticity and integrity of the digital medical images.

Various methods have been widely used to protect digital multimedia files [4, 5]. Among them is digital watermarking that has been proposed and has been also used in the medical field for radiological images [2, 6]. The watermark could be a serial number, like the patient's insurance code, or a hospital logo, or it could include parts or all of the EHR, textual files containing diagnosis, blood test profiles or other biosignals [1, 6, 7]. Moreover, the watermarked image still conforms to the Digital Imaging and Communications in Medicine (DICOM) format.

V. Fotopoulos (✉) · M. L. Stavrinou · A. N. Skodras
Department of Informatics, Computer Science,
School of Science and Technology, Hellenic Open University,
Tsamadou 13–15,
26222 Patras, Greece
e-mail: vfotop1@eap.gr

M. L. Stavrinou
e-mail: marial.stavrinou@gmail.com

A. N. Skodras
e-mail: skodras@eap.gr

M. L. Stavrinou
Department of Physiology, Neurophysiology Unit,
Medical School, University of Patras,
Asclepiou street, Rio Campus,
26500 Patras, Greece

In a medical image during diagnosis, the area of the image that contains the clinical finding constitutes the region of interest (ROI) and it is the one that should be preserved intact. The rest of the part that does not contain any clinical findings is the region of non-interest (RONI) [1–3]. In many works, the watermark is inserted into the RONI, in that way, it leaves the ROI intact without any imperceptible modifications [1–3, 7].

Watermarking can be divided into many categories. Concerning the embedding space, they are categorized into spatial and frequency domain [5]. In the spatial domain, the watermark is inserted by modifying the pixel values or more often the least significant bit [2, 3, 5]. In the frequency domain, the watermark is embedded in some transformed signal coefficients (frequency bands) [1, 8].

In order to retrieve the original image, lately, reversible watermarking (RW) schemes have been introduced. These techniques not only provide protection by embedding the watermark into the original signal but also can recover the original image from the suspected one [9]. Usually, in RW techniques, the RONI is utilized as the embedding area. In such cases, the quality of the RW technique depends highly on the embedding capacity of the RONI. Some authors define RONI as the region of background e.g., the black area inside an X-ray or a magnetic resonance imaging (MRI) image, or as any other nonsignificant area of the image. These methods belong to the *data hiding* category, where the capacity of the carrier is a very important issue.

In the medical field, some newly developed approaches exist for the protection of MRI images. Medical watermarking has been applied in medical MRI images [1, 2, 10]. In order to ensure that the original radiological image can be retrieved in its initial form, reversible watermarking techniques have been introduced [2]. The basic requirements of a reversible data watermarking technique are robustness, imperceptibility, high embedding capacity, readily embedding, and retrieving.

In a work of Coatrieux et al. [10], a mixed reversible scheme was proposed for head MRI images. In the recoverable image tamper proofing approach of Chuang and Chang [8] that employs data compression, the host image was divided into blocks of equal sizes. To recover the “tampered” areas, every block of the host image is compressed by vector quantization to generate the recovery data. In Giakoumaki et al. [7], an area of polygon was chosen to represent the ROI of a medical DICOM image. This ROI remains intact and the watermarking is inserted inside the RONI by difference expansion of adjacent pixel values. In a work by Shih and Wu [11], the ROI is defined as a rectangle inside the original image, which is compressed in a lossless manner and the rest of the image that surrounds the ROI is lossy compressed. The watermark is embedded in the RONI area in the frequency domain.

In the previously described methods, the selection of the RONI is varying. In [1, 3, 7, 11], the ROI is selected manually or automatically as a rectangle. In [10], a simple thresholding method operates in 2×2 pixel blocks to define the ROI with good detail. In [11], the ROI is compressed by lossless compression while the rest, by lossy compression. In our work, a reversible RONI watermarking technique has been implemented for a sequence of brain MRI images. We chose to define ROI first as an area that contains the whole head shape [12]. The ROI itself is the watermark. We consider that it is of importance to preserve the whole brain image for future diagnostic purposes [13]. Moreover, a possible undetected indication in another area not characterized as malicious at the first place could be proven otherwise in a future diagnosis. The embedding capacity is increased by using the total RONI-embedding capacity from all the MRI slice sequences. An automated algorithm detects the available embedding space for each slice and calculates the total embedding capacity. The exact methodology is explained in the [Methods](#) section. The results of the proposed method are given in the [Results and Discussion](#) section, and the conclusions are drawn in the [Conclusions](#) section.

Methods

Overview

Analysis is conducted on a number of MRI slice sequences. MRI scans were acquired using a Siemens Magnetom Vision 1.5 T scanner (Siemens, Enlargen, Germany) according to a magnetization-prepared rapid gradient-echo sequence (256×256 ; field of view, 256; time repetition, 9.7 ms; echo time, 4 ms; flip angle, 12° ; and thickness, 1 mm) at the Institute for Advanced Biomedical Technologies, University G. d’Annunzio, Chieti, Italy. Institutional permission has been granted for the study.

For a single slice as presented in [12] the whole brain shape, the ROI is embedded in the RONI in an adaptive reversible way. An area detection algorithm automatically detects the edges of the smallest rectangle containing the whole head image and counts the number of pixels that belongs to the non-ROI area. Thus, the embedding capacity of this slice is revealed. For improved performance in this work, a new variant of the segmentation algorithm is utilized.

The total RONI area size determines the kind of JPEG2000 compression (lossy or lossless) that will be used for embedding: if the RONI area is big enough to store a losslessly compressed version of the ROI, then the method is fully reversible. If the RONI’s area and capacity are smaller than that, then the ROI image is compressed in the

best possible way, ensuring an optimal image quality. In the latter case, the method is nearly reversible. In both cases, any malicious adversary tampering or bit distortions due to interception are detectable and in such case, the original image is retrieved. The retrieved image is of the best possible quality and moreover, the exact contour of the area of tampering is revealed.

Segmentation

There are many segmentation algorithms in the literature and a pixel-based, thresholding method was initially used [12]. This takes into account the fact that in an MRI slice, there is only one region needed to track the head shape. For segmentation purposes, what's inside or outside is not of importance. The head contour is of interest. For that reason, a simple algorithm scans the image both ways (left to right and right to left) until it reaches a large enough intensity value (that implies that we reached the head area). This is a thresholding method and for 8-bit images, this threshold should be above 15 because we intend to use the four least significant bits of the RONI pixels in the hiding phase. However, for good contrast MRI slices, this threshold may be safely raised up to 60. The largest rectangle containing the ROI is then defined and used as input to the data hiding algorithm. In some literature methods, the ROI is selected by experts and may be a small part of the shape (e.g., a shadow inside the brain) or manually by non-experts. This would greatly assist our method since the capacity of each slice's RONI would be definitely sufficient to hide a smaller area. However, what we want here is a fully automated system, in order for our method to preserve the whole head shape for full recovery. As it was earlier stated, this segmentation method should be kept simple. This may lead to some unwanted phenomena. For example, in Fig. 1a, sample segmentation is shown for one slice. It is observed from the ROI–RONI locations in Fig. 1b (white and black areas, respectively) that for lower thresholds, there may be thin lines coming out of the sides of the head. Such lines are due to noise phenomena and can be a real problem if they are closer to the image edge, comparing to

the real distance of the head from the edge because in that case, the rectangle containing the ROI is unnecessarily large, which consequently reduces the slice's capacity and increases the compressed bitstream's size. Another problem of this method is the fact that even if the ROI is perfectly surrounded by a rectangle, since the head's shape is not rectangular, there are still parts of this rectangle that actually contain RONI pixels.

The ways to solve both of these problems is to use a pixel or block-based algorithm. The second method has been finally selected for this work due to its implementation efficiency. The image is traditionally partitioned into 8×8 pixel blocks. These blocks can be characterized as belonging to ROI or RONI. For this, the average intensity is used as a measure since the RONI blocks' intensity should approach zero. To improve performance, a non-linear, ordered filter is used as a preprocessing stage to remove noise from the imaging device and thus a block characterization map is produced (Fig. 1c). For 256×256 pixels images as those used in this work, a 32×32 binary location map is created that needs to be embedded as additional information (16 bytes) to the bitstream that will be embedded. Since there is always a case that a block can be falsely characterized as ROI, if there is a block, initially tagged as such, that does not have any other ROI blocks around it, it is untagged.

Data Hiding and JPEG2000

For compression, the JPEG2000 was selected both for the efficient lossless compression that it provides and its ability to tweak the compression ratio accordingly, providing best image quality for a given bitrate in lossy compression. However, the main problem in using this method is that the ROI regions that consist of blocks are not rectangular. A solution to this is to rearrange this number of ROI blocks into a rectangle shape (a simplified example is shown in Fig. 2). For this reason, the number of blocks is factorized, and the factors are arranged into two different products in order to produce an array of 8×8 blocks. For example, if the number of ROI blocks is 96, this can be written as $2 \times$

Fig. 1 Segmentation result for slice 24 of person no. 1. **a** Original slice. **b** Segmentation based on the method of [12]. **c** Segmentation based on the proposed method

Fig. 2 A very simple example of ROI blocks reordering to form a rectangle

$2 \times 2 \times 2 \times 2 \times 3$ and if we form two new products by multiplying every second factor starting from the first and the second ones, then we get an array of 8×12 blocks or 64×96 pixels which can be used as input to the JPEG2000 compressor. There are two more issues to be answered about this method. First, what happens if the number of blocks is prime, and second, how efficient is the wavelet-based JPEG2000 compressor if the input image's dimensions are very different (e.g., block array of 2×19 or 33×3). To deal with the first issue, if the number is prime, we add extra 8×8 black blocks. For the second case, we check the ratio $\frac{|N1-N2|}{\max\{N1,N2\}}$, where $N1$ and $N2$ are the dimensions of the 8×8 block array. This value should preferably be 0 (in the case that $N1=N2$), but since this is not always possible, a threshold value of 0.5 has been used in our experiments. Usually, the number of black blocks added is between 0 and 5. Some results of the procedure are given in Fig. 3 where the raw slides and the ROI reordered area (including extra blocks) are shown. It can be noted that there are some black blocks inside the reordered rectangle. These are usually

blocks at the contour of the head which contain part of RONI pixels. To avoid this, a smaller block size can be used (4×4 for example) at the expense of a small increase on the location map. There may also be some noisy blocks that do not really belong to the ROI, but these are very few because both the filtering stage and the second pass (elimination of lonely blocks in the map outside the head contour) remove the majority of these from the location map.

Tamper Detection and Self-Correction

Tamper detection is very straightforward. It can be performed both automatically and manually. In the first case, for each slice, the hidden bitstream is retrieved from the RONI pixels and is decompressed. Then, the retrieved ROI can be separated into small blocks and compared against the current slice view by means of an appropriate threshold value of some parameter (e.g., MSE, variance, etc.). The smaller the size of the blocks, the more accurate the location of tampering becomes. In our case, we dealt with blocks of size 8×8 . Self-correction can be simply accomplished by copying the decoded ROI back into the current view. However, the manual way has a certain advantage. In that case, an expert (e.g., a doctor) is observing both the current slice and the decoded ROI. That makes him/her also capable of identifying the type and probably the reason of tampering (e.g., removing some obvious problems in order to fool an insurance company), something that the automated system can't conclude. Of course, a combined approach would yield the best results; first, the system identifies all tampered cases and second,

Fig. 3 Four different slices of the same MRI sequence and the corresponding reordered ROIs

Table 1 Capacity and compression issues for various subjects

Subject	Slices used	Total bytes needed for lossless compression	Total bytes available through RONI pixels	Fully lossless compression	Average required bitrate (bpp)	Compression required
1	144	4,009,322	2,285,696	No	3.76	53.03%
2	96	1,788,231	2,011,136	Yes	–	–
3	142	3,401,138	2,570,912	No	4.94	38.26%
4	128	3,574,290	2,087,680	No	3.90	51.21%
5	130	3,283,843	2,186,912	No	4.22	47.25%

each one of them is presented to the human expert for consultation.

Results and Discussion

All previous discussion was slice-based. However, MRI is a sequence of images. Thus, it deserves to be seen as a whole and examine better ways to apply the described version in it as if it were a single object. In the current work, the whole MRI sequence is analyzed. Five different sequences were

used, some of them with 111 and some other with 164 slices. The exact number of slices used for each subject is given in Table 1, which also gives the total available capacity of the MRI as well as the total bytes needed for losslessly compressing all slices. For lossless JPEG2000 compression, the Jasper software has been used [14]. There is a case that the total sum of all JPEG2000 compressed files is less than the available capacity, thus, full lossless data hiding is achieved (subject no. 2). In all other cases, lossy compression is required, but the compression required does not lead to significant information losses since it

Fig. 4 Compressed total ROI file size and available capacity per slice for various subjects

varies between 38% and 53%. For compression ratios of 3:1 to 2:1, even lossy compression provides extremely high-quality results. A graph of the required and available bytes per slice and subject are depicted in Fig. 4 where the dashed line corresponds to the first case and the continuous one, to the second. The case of the upper right diagram is notable because it is the only one where all ROIs can be losslessly coded and embedded into the RONIs; this is a case of a female with the smaller head of all the subjects. In contrary, two males with larger heads correspond to the upper left and lower right diagrams where compression of approximately 53% and 51%, respectively, is required. The lower left is a case in between, with a medium-sized head that requires a compression of 38%.

The average required bitrate of Table 1 is an estimate of the best target bitrate that can be obtained if we could afford not to compress in a lossless manner any frame at all. In that case, the whole capacity of the MRI would be equally divided to all slices at the same time. It is also evident that for such a method, there would be a dispersion of the ROI information in different frames. For example, the compressed ROI of frame K would not necessarily be in the RONI of the same frame, but it may occupy the whole or part of frame's L RONI, or even span between the RONIs of frames L and M. For such a scheme, one should have the whole MRI in order to examine and restore a distorted slice. This is one of the major differences between this scheme and that of the slice-based version of [12] or similar.

Another issue has to do with skipping or not some of the first and the last slices which contain less important information. This approach has actually been followed since in all five sequences, the experimenter informed us about how many slices we could skip. If all frames were used, less compression would be required since those frames contain large numbers of RONI pixels. Using all the available capacity of the MRI to losslessly compress the most important slices is also an interesting scenario. The selection of important slices could be human-assisted or automated, by taking the middle slice of the MRI and extending to N slices before and after it.

One last consideration has to do with the number of bits used for embedding. In our previous attempts with the method of [12], both the four and five least significant bits were used. Using five bits was highly beneficial, but in the proposed method, although the capacity would increase by 25%, it is of no meaning, since even the lossy compressed slices are of high quality. In general, one could argue that using four or five bits would produce undesirable noise to the slices. On the other hand, it is also true that for simple viewing, the RONI information is insignificant, thus a simple filter that would zero all the RONI pixels would produce an excellent noise-free view.

One last application scenario would be to use the available capacity in a hybrid way. One could split the capacity in two parts, one simple digital medical data, such as the EHR, and the rest for the self-correction application, of course with a small tradeoff on the average target bitrate.

Conclusions

The role of MRI images for pre-operative assessment and planning, especially for brain neurosurgical as well as for long term follow-up evaluations, is becoming increasingly important. Providing the means for preserving their integrity and authenticity as well as giving self-correction ability is a great aid for the medical community.

The proposed data hiding scheme has been tested under that scope. It is an extension of [12], using MRI sequences as single objects and thus giving new potential to traditional single image/slice data hiding. Five different subjects were used to show differences according to the anatomy of the head. As expected, the size is important since compression could vary from absolutely lossless to near lossless. The scheme performs exceptionally well in terms of available capacity and image quality, and it is flexible in the sense that the capacity can be used in a variety of ways, either automatically, or by the expert's assistance.

Acknowledgments Special thanks to the Institute for Advanced Biomedical Technologies (ITAB), Gabriele D' Annunzio University, Chieti–Pescara, Italy and the Director Prof. Gian-Luca Romani for providing the MRI test sets. Maria L. Stavrinou is currently supported by a State Scholarship Foundation Post-Doctoral grant.

References

1. Navas KA, Thamby SA, and Sasikumar M: EPR hiding in medical images for telemedicine. *Proceedings of the World Academy of Science Engineering and Technology*, Rome, 2008, pp 292–295
2. Coatrieux G, Lecornu L, Sankur B, Roux C: A review of image watermarking applications in healthcare. *IEEE Eng Med Biol Soc* 1:4691–4694, 2006
3. Zain JM, Clarke M: Reversible region of non-interest (RONI) watermarking for authentication of DICOM images. *International Journal of Computer Science and Network Security* 7:19–28, 2007
4. Sencar HT, Ramkumar M, Akansu AN: *Data hiding fundamentals and applications: content security in digital multimedia*. Academic Press Inc, Orlando, 2004
5. Mahdian B, Saic S: A bibliography on blind methods for identifying image forgery. *Signal Processing: Image Communication* 25:389–399, 2010
6. Coatrieux G, Maitre H, Sankur B, Rolland Y, and Collorec R: Relevance of watermarking in medical imaging. *Proceedings of the 3rd Conference Information Technology Application in Biomedicine*, 2000, pp 250–255

7. Giakoumaki A, Pavlopoulos S, Koutsouris D: Secure and efficient health data management through multiple watermarking on medical images. *Med Bio Eng Comput* 44:619–631, 2006
8. Chuang JC, Chang CC: *Medical imaging and augmented reality*. Springer, Berlin, 2004
9. Guo X, Zhuang TG: A region-based lossless watermarking scheme for enhancing security of medical data. *J Digit Imaging* 22:53–64, 2009
10. Coatrieux G, Montagner J, Huang H, and Roux C: Mixed reversible and RONI watermarking for medical image reliability protection. *Proceedings of the 29th International Conference of the IEEE EMBS*, 2007, pp 5653–5656
11. Shih FY, Wu YT: Robust watermarking and compression for medical images based on genetic algorithms. *Information Sciences* 175:200–216, 2005
12. Fotopoulos V, Stavrinou ML, Skodras AN: Medical Image authentication and self-correction through an adaptive reversible watermarking technique. *Proceedings 8th IEEE International Conference on Bioinformatics and BioEngineering*. doi:10.1109/BIBE.2008.4696803, December 8, 2008
13. Connor LT, DeShazo Braby T, Snyder AZ, Lewis C, Blasi V, Corbetta M: Cerebellar activity switches hemispheres with cerebral recovery in aphasia. *Neuropsychologia* 44:171–177, 2006
14. The JasPer project home page, Michael Adams, Available at <http://www.ece.uvic.ca/~mdadams/jasper>. Accessed 3 July 2008.