Anonymization of Burned-in Annotations in Ultrasound Imaging

Mário Antunes, Ricardo Machado, Augusto Silva

Abstract **– Most of the medical imaging modalities compliant with the Digital Imaging and Communication in Medicine (DICOM) standard are currently conveying patient information and other study related information as visual overlay objects on top of the proper imaging contents. Making these image data publicly available for teaching platforms, clinical trials or maintenance support requires some degree of anonymization. In these cases, identity attributes in a given Service-Object Pair (SOP) instance may be simply removed, encrypted or even replaced by dummy values. All these action will take place at the header of the associated DICOM files by changing the proper tag-value data structures and ensuring the overall file integrity. These procedures alone don't ensure full anonymization in some imaging modalities like ultrasound because identifying annotations are sill conveyed in a burned in manner. The anonymization process has to somehow change the pixel data without compromising the underlying imaging contents. Identifying information must then be "blackened" or removed usually by manual or highly user driven methods. This approach is neither systematic nor feasible for large scale deployment of image databases. This paper proposes a smart approach to the problem of anonymization of burned-in annotations in ultrasound imaging studies. Our approach is based on customized character template matching. A character template database is firstly built during a learning phase using a representative sample of ultrasound studies produced by different manufacturers. During the matching the phase, using the proper identifying attributes from the header of the DICOM files, bit-mapped string templates are arranged and then matched against the whole pixel data. This approach has proven to be robust and yields very high matching scores providing a solution capable of effectively de-identifying burned in annotations in ultra-sound images independently of their position within the image.**

Keywords **– Medical Imaging, Ultrasound, DICOM, Anonymization, Burned in Annotations, Image Databases**

I. INTRODUCTION

Large scale Medical Image Databases with several access profiles or totally public access impose privacy protection. The raw material that embodies these databases often come from Picture Archiving and Communication Systems (PACS) housed by healthcare institutions. Images are made persistent objects when they become stored as files and sent to the archive for further reviewing or processing. A whole set of services and protocols are required to accomplish the different imaging workflows that may occur within an Imaging Department. These services and protocols are currently governed by the DICOM standard. The instantiation of a service over an object image is called in DICOM terminology a SOP instance. Sending an image to an archive server is an example of a frequent SOP instance. The DICOM standard uses a tag-value-length (TVL) file layout requiring many mandatory attributes besides the proper pixel data field. Some of these are related with the specifics of the imaging study while others are patient related.

Most of the medical imaging modalities are currently conveying patient information and other study related information as graphical overlays on top of the proper imaging contents. Making these image data publicly available for teaching platforms, clinical trials or maintenance support requires some degree of anonymization in order to ensure privacy. In these cases, identity attributes in a given SOP instance may be simply removed, encrypted or even replaced by dummy values. All these actions take place at the meta-data fields of the associated DICOM files by changing the proper tag-value data structures and ensuring the overall file integrity. These procedures cannot be applied to some imaging modalities like ultrasound because identifying annotations are sill conveyed in a burned in manner. The anonymization process has to somehow change the pixel data without compromising the underlying imaging contents. Identifying information must then be "blackened" or removed usually by manual or highly user driven methods. This approach is neither systematic nor feasible for large scale deployment of image databases. The most obvious case is concealing the patient name as shown in the example of figure 1

This paper proposes a smart approach to the problem of anonymization of burned-in annotations in ultrasound imaging studies. Our approach is based on customized character template matching. A character template database is firstly built during a learning phase using a representative sample of ultrasound studies produced by different manufacturers. During the matching phase, using the proper identifying attributes from the header of the DICOM files, bit-mapped string templates are arranged and then matched against the whole pixel data. High matching scores will point to pixels that should be masked concealing thus potential identifying. Notice that this a irreversible process since the original pixel data cannot be recovered after this transformation. For the purpose of maintaining DICOM compliance a new image is created under a new SOP instance with its own unique identifier (UID). Section II of this paper describes the methods and materials that were proposed and used in this anonymization task, Section III highlights and discusses some relevant results and Section IV concludes and presents some prospects for future work.

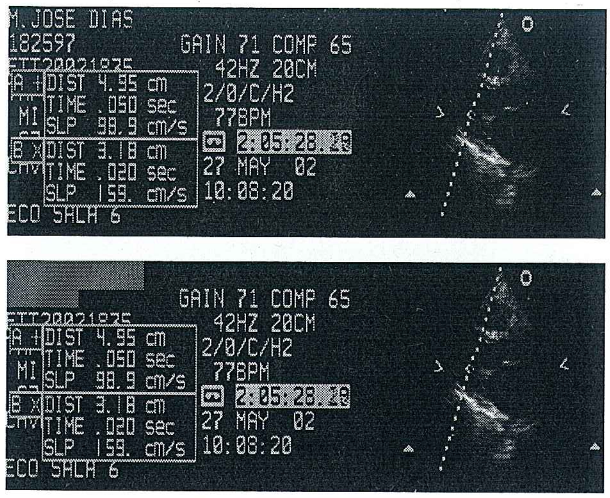

Fig. 1 - **Concealing private information. Top:** region of a medical image which contains private information. **Bottom:** the same region with the private information concealed
Note: Patient name is a pseudonym.

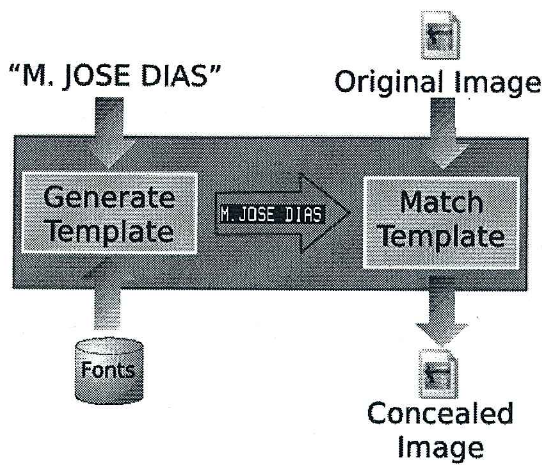

Fig. 2 - Architecture of the technique proposed in this paper.

II. Materials and Methods

As mentioned before our method is conceptually grounded on template matching techniques. The novelty here comes from the fact that the templates are bit-mapped versions of potential identifying strings that may be readily extracted from the DICOM meta-data. In our case the most important fields that underwent concealment processing were the patient name and study code. As shown in figure 2, our anonymization process has as input a set of identifying strings and a target image where bit-mapped versions of these strings appear and should afterwards be somehow concealed.

When confronted with a generic optical character recognition (OCR) approach followed by string matching [1], our strategy is definitely more focused and efficient regarding the computational burden. The lack of generality is overcome by having a sufficiently broad character database that is used to build candidate templates based on the string information required by the anonymization rules.

A. Template Generation

Two main techniques may be used to create a string based template. A very rapid manner would consist on using the host system fonts to create the template if one of these fonts would have an acceptable level of similarity with the fonts used by medical ultrasound equipment to generate bit-mapped annotations. However, since the fonts used in medical images are not known in advance and ultrasound machines use a very limited if not unique font set it is not possible to use this first technique. Our technical option focused on a fairly more generic and robust alternative that relied on the concatenation of individual character templates.

The first step in using the second technique is to create a representative database of individual characters that are likely to appear in the images. To achieve this necessary goal we gathered a large number of ultrasound images produced by a significant collection of equipments from different manufacturers. After a tedious manual cut and paste process single character binary templates were generated and stored in a database. This preliminary but necessary task may be called a training phase since all the sought for character sets from several font types were in this way embedded in the anonymization system. As expected, the performance of this technique heavily depends on this phase since the quality of the individual templates will determine the effectiveness of the overall anonymization process.

B. Template Matching

For each string to be concealed and for each font type a set of templates is generated by simple concatenation of the character masks that compose the string. It is worth mentioning that our solution is able to deal both with monospaced and proportional character sets. The burned-in annotations are detected and thereafter concealed by matching the candidate templates throughout the image by computing local similarity metrics such as the one given by simple correlation [2]:

$$R(x,y) = \sum_{x',y'} T(x',y')I(x+x',y+y') \qquad (1)$$

where $T(x',y')$ is the template window centered at coordinates (x,y) in the input image $I(x,y)$ space.

For the sake of generality no assumptions are made about preferred locations of the burned-in annotations. Practice has shown that these locations may differ among the several types of ultrasound equipment that were tested. Each image is thus completely scanned in order to ensure that all possible string locations are detected. A 100% correlation score indicates a perfect match and the rectangular region underneath the template window may be safely masked. However, as shown in figure 3, due to image noise and the specifics of the template generation process, most of the times pixel wide misalignments occur reducing the matching scores. Heuristically we decided to mask all regions exhibiting a correlation score higher than 85% for totally explicit patient names and 95% for abbreviated patients names with a "." character. Due to the nature of the correlation operation and to the possibility of having multiple instances of

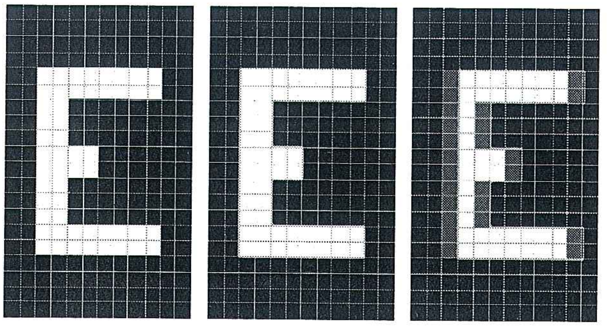

Fig. 3 - **Matching examples. Left:** Template of character 'E'. **Middle:** correctly positioned template with 100% similarity. **Right:** Template shifted 1 pixel with 90% similarity.

burned-in annotations the accurate location of the regions to be masked was accomplished by a local maximum search over the correlation output image. It was not necessary to perform any kind of thresholding before matching the image against the binary character templates since bit-mapped annotations always appear as saturated pixels in a gray-level representation. Another important prior assumption is that annotations don't scale even if image content is subject to some form of zooming.

III. RESULTS

After building the character database we tested our solution over several hundreds of cardiac ultrasound studies. Anonymization was carried out both for static images and video. In this case the whole algorithm was applied on frame by frame basis. In fact this is a rather brute force approach since burned-in annotations usually don't change position in multi-frame object. There are however exceptions where the clinical viewing mode may change in the middle of an acquisition or even the display itself may also be altered to split screen layouts in a manner that identifying annotations escape from their initial locations.

To illustrate our results we present several figures where the performance of the concealment process may be visually assessed. Common identifying DICOM tags such as Patient Name, Patient ID, Dates, Accession Number, Institution ID were extracted from the image files meta-data and fed into the anonymization workflow.

Figures 4 and 5 show the most common situation with B-Mode images where annotations are usually placed outside the triangular field of view. All the masked regions correspond to pixels that were part of bit-mapped written information associated with the prior list of tags for anonymization. Sometimes, frames appear with pop-up measurement windows over the accession number (or other sought for bit-mapped string). In these cases the recognition task fails for almost full coverage as is the case shown in figure 6. However anonymization here is in practice implicit as the accession number is hardly visible.

In Figure 7 a B-mode study from a different equipment is shown. All annotations are located on the top ribbon of the image and are completely masked.

M-mode images are also frequent in cardiac ultrasound studies. Figures 8 and 9 show M-mode images with the required ID tags appropriately masked.

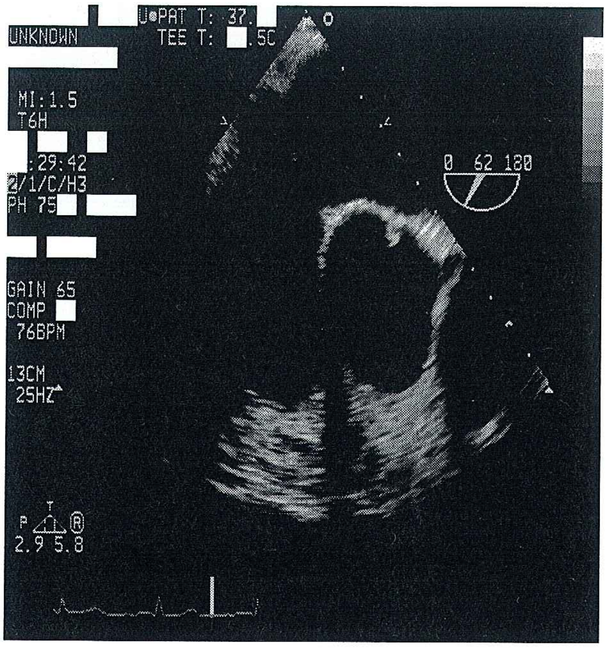

Fig. 4 - Standard cardiac B-mode ultrasound image with burned-in information masked

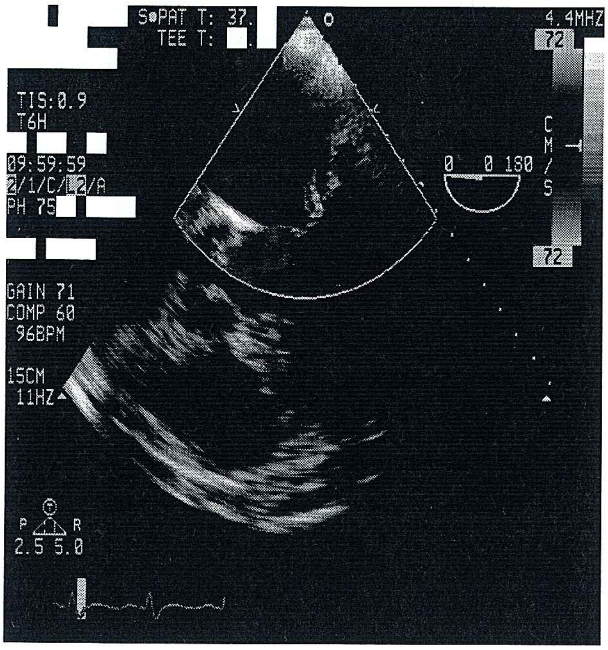

Fig. 5 - Standard cardiac B-mode ultrasound image with Doppler color overlay and burned-in information masked.

All the tests performed over a repository comprising hundreds of cardiac ultrasound studies indicate a very high success rate regarding the proposed anonymization workflow. It is evident that this success rate is highly dependent on the quality of the prior character set database that must be built from representative training samples taken from the repository. As it can be seen in figures 7 to 9, there are cases where noise within the regions comprising annotations to be masked may impair a correct result due to poor matching.

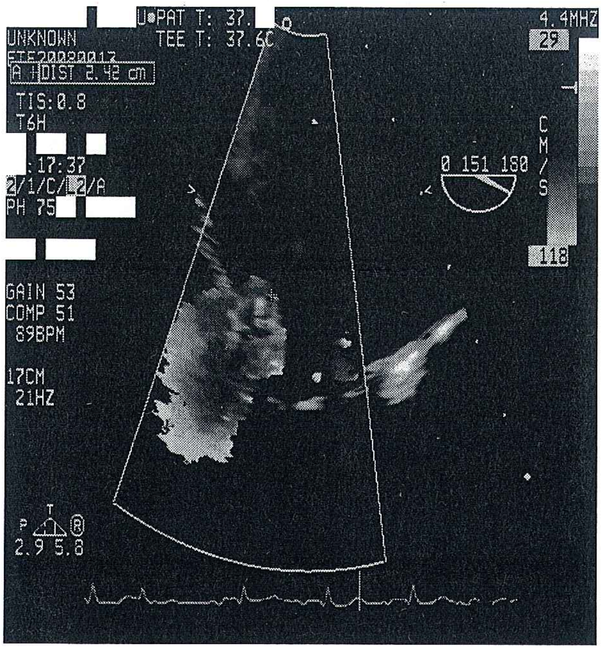

Fig. 6 - Pop-up window almost covers the accession number and local matching fails.

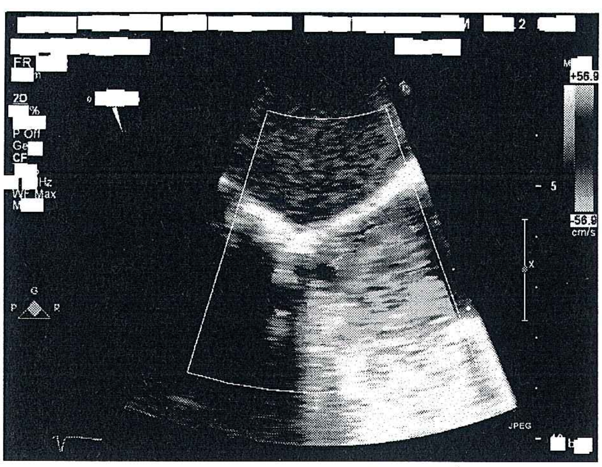

Fig. 7 - Standard cardiac ultrasound image produced by an alternative equipment.

Fine tuning the thresholding levels may solve this problem at the expense of over segmentation of smaller pieces of similar annotations although without compromising neither the anonymization task nor the proper image content.

IV. CONCLUSION

In this paper, we have presented a technique to conceal private information in medical ultrasound images where the overall anonymization procedure must also handle the masking of burned-in identifying annotations. The results show that this technique is effective in identifying and concealing different types of information in various types of medical images with burned-in annotations provided there is representative character-set database. From this database character and string templates may be formed driven by DICOM ID tags contained in the image file metadata. Template matching by simple correlation proved to be effective to detect the candidate burned-in annotations. The quality of the prior character-set database and noise in the burned-in annotations are the principal factors determining the success rate of this method.

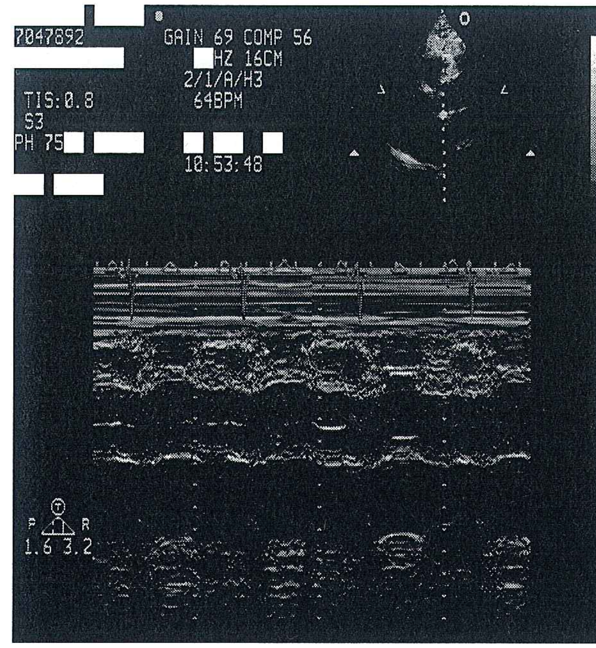

Fig. 8 - Cardiac M-mode ultrasound image. Some ID tags are located in different positions when compared to the B-mode standard view

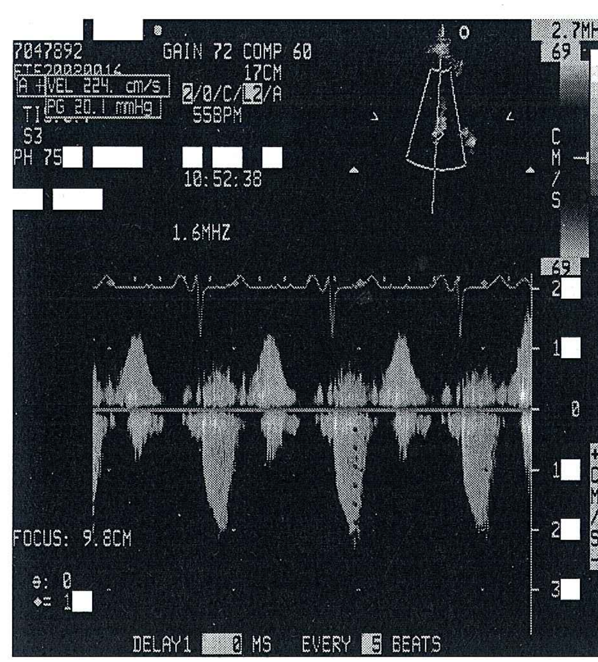

Fig. 9 - The same as before with a pop-up window over the accession number.

REFERENCES

[1] Shunji Mori, Hirobumi Nishida, and Hiromitsu Yamada, *Optical Character Recognition*, John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1999.

[2] Roberto Brunelli, *Template Matching Techniques in Computer Vision: Theory and Practice*, Wiley Publishing, 2009.

[3] NEMA, "Dicom", `http://medical.nema.org/`, fev. 2011.