

Preserving Privacy in Structural Neuroimages

Nakeisha Schimke, Mary Kuehler, and John Hale

Institute of Bioinformatics and Computational Biology
The University of Tulsa
800 South Tucker Drive
Tulsa, Oklahoma 74104

Abstract. Evolving technology has enabled large-scale collaboration for neuroimaging data. For high resolution structural neuroimages, these data are inherently identifiable and must be given the same privacy considerations as facial photographs. To preserve privacy, identifiable metadata should be removed or replaced, and the voxel data de-identified to remove facial features by applying skull stripping or a defacing algorithm. The Quickshear Defacing method uses a convex hull to identify a plane that divides the volume into two parts, one containing facial features and another the brain volume, and removes the voxels on the facial features side. This method is an effective alternative to existing solutions and can provide reductions in running time.

Keywords: Medical image privacy, neuroimaging, de-identification, HIPAA.

1 Introduction

The digitization of health records and medical images has transformed healthcare and medical research. New technologies provide instant access to patient and subject data by automatically disseminating the information to healthcare providers and research collaborators. Expanded storage and transfer capabilities have made feasible the addition of medical images to these electronic records, but as the demand for capturing and storing images increases, so does the need for privacy measures. For shared data sets, the need for removing protected health information (PHI) is agreed upon, but the extent to which medical images constitute PHI is still debated.

The Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule [12] defines “full face photographic images and any comparable images” as PHI. With respect to identifiability, high resolution structural magnetic resonance imaging (MRI) datasets are comparable to full face photographs, and volume rendering software is freely available. Fig. 1 is a volume rendering of a structural MRI using 3D Slicer [1], an open source software package for medical image analysis. The result is clearly identifiable as a human face.

The challenges of removing identifiable metadata are well documented, and there are numerous tools for automating the process. There are also formal



Fig. 1. Volume rendering using 3D Slicer. Sample MRI data is from 3D Slicer.

models for privacy, such as *k-anonymity* [23]. However, the inherent privacy risks of the neuroimages themselves is less well defined. The relative anonymity of subjects in structural MRI may be compromised by the image itself. This paper explores the potential privacy hazards associated with neuroimage datasets. It also proposes a new algorithm for image-based de-identification of neuroimages and evaluates its effectiveness and performance.

2 Background

Large scale collaborative research efforts have the potential to transform neuroscience. The Alzheimer’s Disease Neuroimaging Initiative (ADNI) [2] is a multisite collaborative research project that has collected images from over 40 sites and distributed data to more than 1,300 investigators to date [13, 15]. Its success has inspired similar initiatives for other diseases.

There are, however, obstacles to neuroimage data sharing that hamper collaboration. Solutions to technical challenges, including data storage, transmission, management, and dissemination, continue to evolve. The task of maintaining subject privacy while disseminating data has made significant progress. Metadata removal is routinely integrated into the scientific workflow. However, the determination of when and how to apply de-identification to the neuroimage itself has yet to be made. The benefits of sharing neuroimaging data are clear, but pressing concerns over subject privacy must first be addressed.

The terms *anonymization* and *de-identification* are often used interchangeably, but their subtle differences are significant to subject privacy. The core idea mechanism for patient privacy relies on obscuring the subject’s identity by hiding medical and personal data, often applied to meet the de-identification requirements of HIPAA. A dataset de-identified under the HIPAA Privacy Rule can be distributed and used. HIPAA designates eighteen identifiers as PHI, including “full face photographic images and any comparable images” [12].

Anonymization is not as clearly defined. True anonymity would prevent a dataset from ever being re-identified but is difficult to achieve while retaining useful data [18]. Neuroimaging studies often require metadata such as gender and age for analysis, and removing these could negatively impact results. Practical

anonymity inhabits a grey area between true anonymity and an acceptable yet undefined limit to the possibility of re-identification.

The need for re-evaluation of PHI is evident when a few pieces of seemingly innocuous data can be re-linked to identify a subject. Medical images belong to a class of health data that is inherently self identifying and laden with contextual information about the subject, their condition, treatment, and medical and personal history. The privacy issues associated with the storage and use of medical images warrant special consideration, and the current approaches of simply removing metadata may be insufficient.

3 Privacy Issues in Medical Images

While textual data can be redacted by simply removing or replacing the offending field, the image, which can constitute self identifying data, is not so easily sanitized. Removing identifying features in medical images may destroy the very information a researcher needs.

Table 1. Threats to subject privacy from medical images.

Type	Description	Example
Direct	Reveals a condition	X-ray reveals fractured wrist
Re-linkage	Metadata reveals identity	Metadata includes gender, age, and zip code and tied back to patient
Existential Inference	Image known to exist	Subject in imaging study assumed to be a case rather than control
Identification	Inherently identifiable	Facial features identify subject

The primary threats to subject privacy from medical images are listed in Table 1. A direct threat occurs when the image reveals a condition or other private information, but a more likely scenario is re-linkage, where the image is used to identify the subject along with metadata. The existence of a medical image or participation in a study may also suggest the presence of a condition, perhaps incorrectly. Neuroimages are particularly challenging because they are inherently identifiable. High resolution neuroimages contain detailed facial features that can be used to re-identify the subject. The neuroimage could be used to discover an identity from a large database of faces or to confirm a subject's identity.

3.1 Neuroimage Re-identification

There are many potential avenues for re-identifying a subject using their neuroimage. Re-identification occurs in two phases, reconstruction and recognition.

The reconstruction phase produces a likeness of the subject to be used in the recognition phase for discovering the subject's identity.

In forensic science, facial reconstruction requires a blend of artistic and scientific skills to reproduce a likeness of the subject. Reconstruction is more straightforward using structural MRI because of the high spatial resolution. Several packages for analyzing neuroimage data provide built-in volume rendering capabilities, including AFNI [3], 3D Slicer [1], and MRICron [20]. Typical volume rendering software offers the ability to change lighting conditions and viewing angles. These features can be used to match rendered volumes against photographic facial images.

Facial recognition can be applied using a variety of techniques to achieve novel identification, attempting to discover an identity, or identity confirmation. Metadata can be used to guide a facial recognition search, narrowing down the potential subjects using basic non-PHI fields such as gender and age. The current limitations and relatively poor performance of facial recognition techniques make it tempting to dismiss the potential for re-identification based on flawed assumptions: (1) facial recognition will never improve, and (2) only correct identifications are problematic. The latter fails to consider the damage caused by incorrect identification. Challenging a false re-identification may require the individual to reveal their records.

The problems plaguing facial recognition techniques are not easily confronted, but researchers in the field are making progress. Facial recognition techniques are detailed with links to recent advances at the Face Recognition Homepage [10]. A NIST report on face recognition illustrates significant improvements in the field[11].

Hardware advances can also improve the results of facial recognition. Increased storage capacity and computing power allow higher quality images to be stored and compared more quickly. Facial recognition software struggles when viewing angles and lighting vary [24], but volume rendering software can generate multiple images with a wide range of light sources and angles to match source photographs. Therefore, if neuroimage-based recognition can perform with comparable results, they must be offered the same protection.

3.2 Neuroimage De-identification

There are two common approaches to de-identifying neuroimages, skull stripping and defacing. Skull stripping is the identification and removal of non-brain tissue as part of the typical analysis workflow. It has many benefits, including improved registration between images, removal of acquisition artifacts [22], and de-identification by removing facial features.

There are several methods for skull stripping, and many are integrated with widely used neuroimage analysis software [3, 5, 9, 22, 8, 21]. Several skull stripping methods are compared and analyzed in detail in [7]. Skull stripping methods are highly sensitive to parameters, which may often result into loss of desirable brain tissue. The results may also vary between methods and can require manual

correction. Differences in data sets may impact further analysis, such as segmentation. Skull stripping may also favor a particular region based on the particular study [6]. This complicates meta-analysis, data re-use, and collaboration by discarding potentially relevant voxels.

Unlike skull stripping, defacing techniques [6] preserve non-brain tissue. The MRI Defacer approach removes only voxels with zero probability of containing brain tissue and non-zero probability of containing facial features using a manually labeled face atlas. The result appears as though the facial features were eroded, leaving the brain volume intact.

It is tempting to de-identify with skull stripping since it is part of analysis, but defacing techniques allow for more flexibility. Simply skull stripping an image may discard useful data. Defacing is an effective method for removing facial features, and it does not interfere with subsequent analysis. MRI Defacer relies on a face atlas to identify features, which may not apply well to all datasets.

4 Quickshear Defacing

Quickshear Defacing is a new technique for removing facial features from structural MRI. The primary objective is to provide an efficient and effective defacing mechanism that does not rely on external atlases. It uses a binary mask to identify the brain area to protect, as illustrated in Fig. 2. It identifies a plane that divides the volume into two parts: one containing the brain volume and another containing facial features. The voxels that fall into the latter volume \mathbf{F} are removed, leaving the brain volume \mathbf{B} untouched. Removing all facial features is not necessary to de-identify the image, and the subject’s identity can sufficiently be obscured by removing the primary features (eyes, nose, mouth).



Fig. 2. Quickshear Defacing illustrated (left). Sample slice (middle) and volume rendering (right) after defacing.

The brain mask is created using a skull stripping technique, with the flexibility to use an existing skull stripped volume. Non-brain tissues such as cerebrospinal fluid and the optic nerve, among others, are often problematic for skull stripping techniques, which aim to include only brain tissue. Quickshear, however, does not need to fully distinguish between brain and non-brain tissue. To

reduce complexity and simplify the process, a flattened, two-dimensional sagittal view of the brain is considered. The edge mask is used to find the convex hull. By definition, the convex hull of the brain will form a polygon so that all brain voxels are either on the boundary or inside.

Andrew’s monotone chain algorithm is used to find the convex hull [4], The algorithm sorts the points lexicographically and finds the lower and upper halves of the hull. Selecting the leftmost point (x_0, y_0) ¹ and the adjacent point (x_1, y_1) on the hull ensures that all of the brain voxels are contained in the remaining portion of the hull.

The three-dimensional defacing mask is created by discarding all voxels that lie below the line formed by the points defined by

$$w_j = \left(\frac{y_1 - y_0}{x_1 - x_0} \right) (j - x_0) + y_0 - b . \quad (1)$$

The value of b specifies a buffer to ensure preservation of the brain volume by shifting the line by $-b$ values in the j direction.

The methods were tested with the Multimodal Reproducibility Study data set from Landman, et al., using MPRAGE scans with a $1.0 \times 1.0 \times 1.2$ mm³ resolution. Acquisition is detailed in [14]. The data set contains 42 images from 21 health subjects. Defacing was performed on Ubuntu 10.10 running in VirtualBox on an Intel i7-2600k with 2GB RAM. Running time is shown in Table 2 as an average per image, averaged over five runs.

Table 2. Performance for defacing per image of sample data set, averaged over five runs.

Method	Skull Stripping Time (s)	Defacing Time (s)
MRI Defacer	-	260.17
Quickshear	205.71	4.30

By design, Quickshear Defacing should not remove any voxels identified as brain by the binary mask it is given. This is a basic sanity check, where the defaced volume is compared voxelwise with the brain mask identified by each of three skull stripping techniques (AFNI 3dSkullStrip, FSL BET, and FreeSurfer HWA). On average, Quickshear Defacing discarded fewer brain voxels from fewer images than MRI Defacer.

Volume rendering was applied using MRICron [20] to the resulting defaced images and passed through the OpenCV Haar classifier [19] to detect faces. For Quickshear, 12 of 42 images were classified as containing a face, and for MRI Defacer, 9 of 12 contained faces. Quickshear tended to leave behind features such

¹ The leftmost point is chosen as the starting point based on a space where $+x$ -axis is the inferior to superior (front to back).

Table 3. Average number of brain voxels discarded for each defacing mechanism (Number of images with voxels discarded).

Defacing Method	Brain Mask		
	AFNI	BET	HWA
MRI Defacer	408.74 (12)	75271.93 (42)	422.0 (7)
Quickshear	0.0 (0)	5560.76 (13)	0.0 (0)

as the eye sockets and nasal cavity that may be triggering a false positive. Upon visual inspection, defacing appeared adequate using both methods. MRI Defacer left behind extreme features like the nose in some cases.

5 Conclusions

While the practical and effective discussion concerning privacy in structural neuroimages continues, there are effective measures that can be taken immediately to improve subject privacy. Adopting such measures to protect both metadata and pixel data can increase the flow of data both internal and external to research organizations and encourage collaboration.

Metadata can be removed using existing anonymizing tools, such as the LONI De-identification Debabelet [16] and DICOMBrowser [17]. To remove pixel data, skull stripping or one of the defacing algorithms is recommended. Skull stripping is an effective method for removing facial features, but it may discard desirable tissue. If reproducibility and peer review are the motivations for data sharing, skull stripping may be sufficient and can save time if it is part of the workflow. For data reuse, a defacing approach such as the one presented in this paper may be preferred.

Quickshear Defacing uses a two-dimensional view of the data to create a convex hull, which identifies a plane that divides the volume into two parts, one containing the entire brain and the other facial features. By removing all voxels on the face side, the image data is de-identified.

Quickshear Defacing preserves more brain voxels in more images than MRI Defacer. After MRI Defacer, fewer volumes were identified as containing faces by the Haar classifier. Visual inspection of both techniques showed that the remaining volumes were unlikely to be identified.

Further tests on the data should be applied to determine the effects of the new defacing technique proposed in this paper on further skull stripping. Additionally, implementing other techniques in addition to the Haar classifier to verify the removal of facial features may illuminate the performance of both defacing methods.

Acknowledgments

We gratefully acknowledge support from the William K. Warren Foundation.

References

- [1] 3D Slicer. <http://www.slicer.org> (Accessed 2010)
- [2] ADNI: Alzheimer's Disease Neuroimaging Initiative. <http://www.loni.ucla.edu/ADNI/> (Accessed 2011)
- [3] AFNI. <http://afni.nimh.nih.gov> (Accessed 2011)
- [4] Andrew, A.M.: Another efficient algorithm for convex hulls in two dimensions. *Inform Process Lett* pp. 216–219 (1979)
- [5] BET - Brain Extraction Tool. <http://www.fmrib.ox.ac.uk/fsl/bet2/index.html> (Accessed 2010)
- [6] Bischoff-Grethe, A., et al.: A technique for the deidentification of structural brain MR images. *Hum Brain Mapp* 28, 892–903 (2007)
- [7] Fennema-Notestine, C., et al.: Quantitative evaluation of automated skull-stripping methods applied to contemporary and legacy images: Effects of diagnosis, bias correction, and slice location. *Hum Brain Mapp* 27, 99–113 (2006)
- [8] FreeSurfer. <http://surfer.nmr.mgh.harvard.edu> (Accessed 2011)
- [9] FMRIB Software Library. <http://www.fmrib.ox.ac.uk/fsl/> (Accessed 2010)
- [10] Grgic, M., Delac, K.: Face recognition homepage. <http://face-rec.org> (Accessed 2011)
- [11] Grother, P.J., Quinn, G.W., Phillips, P.J.: Report on the evaluation of 2D still-image face recognition algorithms. Tech. rep., National Institute of Standards and Technology (2010)
- [12] HIPAA Administrative Simplification: Regulation Text (2009)
- [13] Kolata, G.: Rare Sharing of Data Leads to Progress on Alzheimers. *New York Times* (August 12 2010)
- [14] Landman, B.A., et al.: Multi-parametric neuroimaging reproducibility: A 3T resource study. *NeuroImage* (2010)
- [15] Mueller, S.G., et al.: Ways toward an early diagnosis in Alzheimer's disease: The Alzheimer's Disease Neuroimaging Initiative (ADNI). *Neuroimag Clin N Am* (2005)
- [16] Neu, S.C., Valentino, D.J., Toga, A.W.: The LONI Debabeler: a mediator for neuroimaging software. *NeuroImage* 24, 1170–1179 (2005)
- [17] Neuroinformatics Research Group: DICOM Browser. <http://nrg.wustl.edu/software/dicom-browser/> (Accessed 2011), Washington University School of Medicine
- [18] Ohm, P.: Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Rev* 57(6) (2010)
- [19] OpenCV. <http://opencv.willowgarage.com> (Accessed 2011)
- [20] Rorden, C.: MRIcron. <http://www.cabiatl.com/mricro/mricron/index.html> (Accessed 2010)
- [21] Ségonne, F., et al.: A hybrid approach to the skull stripping problem in MRI. *NeuroImage* 22, 1060–1075 (2004)
- [22] Smith, S.M.: Fast robust automated brain extraction. *Hum Brain Mapp* 17, 143–155 (2002)
- [23] Sweeney, L.: k-anonymity: a model for protecting privacy. *Int J on Uncertain Fuzz* 10(5), 557–570 (2002)
- [24] Zhao, W., Chellappa, R., Phillips, P.J., Rosenfeld, A.: Face recognition: A literature survey. *ACM Comput Surv* 35, 399–458 (2003)