

Privacy Protection in MRI Scans Using 3D Masked Autoencoders

Lennart A. Van der Goten^{1,2}(✉) and Kevin Smith^{1,2}
*for the Alzheimer's Disease Neuroimaging Initiative***

¹ KTH Royal Institute of Technology, Stockholm, Sweden
{lavdg,ksmith}@kth.se

² SciLifeLab, Solna, Sweden

Abstract. MRI scans provide valuable medical information, however they also contain sensitive and personally identifiable information that needs to be protected. Whereas MRI metadata is easily sanitized, MRI image data is a privacy risk because it contains information to render highly-realistic 3D visualizations of a patient's head, enabling malicious actors to possibly identify the subject by cross-referencing a database. Data anonymization and de-identification is concerned with ensuring the privacy and confidentiality of individuals' personal information. Traditional MRI de-identification methods remove privacy-sensitive parts (e.g., eyes, nose etc.) from a given scan. This comes at the expense of introducing a domain shift that can throw off downstream analyses. In this work, we propose CP-MAE, a model that de-identifies the face by remodeling it (e.g., changing the face) rather than by removing parts using masked autoencoders. CP-MAE outperforms all previous approaches in terms of downstream task performance as well as de-identification. With our method we are able to synthesize high-fidelity scans of resolution up to 256^3 on the ADNI and OASIS-3 datasets – compared to 128^3 with previous approaches – which constitutes an eight-fold increase in the number of voxels.

Keywords: Magnetic Resonance Imaging · Privacy in Health · Generative Modeling.

1 Introduction

While MRI scans are usually visualized as 2D slices, it is also possible to render a high-quality 3D model using techniques such as volumetric raytracing, yielding a realistic depiction of a patient's face. This is problematic in terms of data

** Data used in preparation of this article were obtained from ADNI (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf




Fig. 1: MRI scans pose a privacy risk since highly-realistic face renderings can be crafted and misused for malicious purposes. The *remodeling* approach to MRI de-identification, used in our work, retains the brain (=) and remodels all other features ($\neq$). (*top*: 3D view, *bottom* slice-view, taken from OASIS-3)

privacy: Given an MRI scan and a face database with associated identities, a malicious actor can find the closest match to a given face rendering, allowing them to *potentially infer* a patient’s identity.

Various de-identification methods have appeared over the years that address the removal of privacy-sensitive parts. Most are quite simple and differ in their level of aggressiveness: While FACE MASK [10] merely blurs out the face, skull-stripping methods such as MRI WATERSHED [17] retain only the brain and remove everything else. However, these traditional approaches are potentially problematic. Many tools in the medical workflow require the presence of certain landmarks as preconditions. If those landmarks are absent, the analysis might be impaired or, in the worst case, it might be impossible to perform the analysis at all [18]. On the other hand, if a method is not aggressive enough it might still be feasible to infer a patient’s identity, a fact that proves especially harmful when such a tool is to be used automatically and without oversight on a large corpus of scans.

Our recently proposed GAN-based approach named CP-GAN [5] introduced a new paradigm: to remodel the face and leave the medically-relevant information intact. Simply put, a patient’s MRI scan is de-identified by copying their brain and remodeling the remaining parts such that they appear to be an actual MRI scan, but do not give away information about the patient’s real face or identity. This approach eliminates the aforementioned trade-off, as it ensures reasonable placement of landmarks and thus minimizes domain shift, leading to a safer and more effective de-identification.

Recently, the emergence of *masked autoencoders* (MAEs) has propelled generative approaches forward, departing from the previously dominant GAN method-

ology by predicting stochastically-masked segments of the data instead of formulating an adversarial game. MAEs offer two fundamental advantages over GANs: (i) Higher training stability, and, (ii) Higher data efficiency. Considering that MRI de-identification with generative models involves high memory requirements, data scarcity, and often suffers from instability due to small batch sizes, MAEs appear to be an attractive alternative.

In this work, we adapt masked autoencoders for MRI de-identification. Our contributions are as follows:

1. We propose CP-MAE, a *masked autoencoder*-based model that can deal with high-resolution MRI scans. Whereas previous remodeling-based methods were limited to a resolution of 128^3 , CP-GAN can produce de-identified scans of size 256^3 , an 8-fold increase in the number of voxels.
2. To the best of our knowledge, CP-GAN is the first to combine a volumetric VQ-VAE with an MAE for 3D MR image synthesis
3. We demonstrate that CP-MAE yields superior de-identification performance compared to other methods and can be robustly trained on modestly-sized datasets (e.g., ADNI, OASIS-3)
4. We show that de-identification with CP-MAE introduces minimal adverse effects on brain tissue and subcortical segmentation tasks compared to other approaches.

2 Related Work

MRI De-Identification. MRI de-identification is a critical pre-processing step in neuroimaging, and many methods have been proposed to effectively perform this task. BET [19] is widely used for its simplicity. Despite its effectiveness, it sometimes fails to exclude non-brain tissues, such as dura and eyes. ROBEX [7] aims to address some of these shortcomings by employing a trained random forest classifier to distinguish brain and non-brain voxels. MRI WATERSHED, in turn, [17] employs a watershed transformation followed by a deformable model to extract the brain. Milchenko et al. [10] propose FACE MASK, focusing on the protection of sensitive facial information in MRI data. This method combines the benefits of defacing and skull-stripping by generating a mask that can be used to blur out facial features while preserving brain anatomy. Schimke et al. introduced a tool named QUICKSHEAR [16], which is intended to compute a hyperplane in 3D space that separates the facial features from the brain. DEFACE [1] is a deformable model that estimates which voxels belong to the brain, said voxels can then be cut out for de-identification purposes.

Finally, CP-GAN [5] remodels non-brain tissues as opposed to removing them. The de-identification is performed by a conditional GAN that takes as input the patient’s brain and a convex hull hinting at the metric extent of the skull to be generated. Our approach follows this paradigm of remodeling the face and skull, as opposed to opting for something that would remove these parts altogether.

Image Synthesis. Generative models, including VAEs [8], GANs [4], diffusion models [20], and masked autoencoders (MAEs) [3,13,14], are extensively studied.




Fig. 2: **Vector Quantization Stage.** We preprocess the data in two independent stages: (i) We extract the brain, and simultaneously (ii) remove the brain from the full skull. Both representations are then *compressed* by two independent VQ-VAEs into discrete 3D volumes.

MAEs excel by reconstructing an image from a masked version of the image. Our work builds upon recent MAEs like MaskGIT [2] and Paella [15].

3 Method

Task Definition. As in [5], given a set of 3D scans $(X^{(i)})_{i=1,\dots,N}$ following a data distribution $\mathcal{P}_X$ and having a resolution of S^3 , we seek a mapping $Y = G_\Phi(\gamma(X))$ parameterized by Φ that de-identifies a *raw* scan X and yields Y . The purpose of the *privacy transform* $\gamma(\cdot)$ is to provide $G_\Phi(\cdot)$ with a minimal blueprint to guide the skull synthesis without leaking information about the patient’s facial features³, thus $\gamma(X)$ should at least include a representation of the brain to inform the model about the metric constraints of the to-be-synthesized skull.

Overview. Synthesizing (3D) MRI scans is challenging in terms of memory as a volume contains a cubic number of voxels. VQ-VAE is therefore particularly attractive as the number of voxels could in the most extreme case be reduced from 256^3 to a mere 64^3 , reducing the overall number of voxels by a factor of 64. Our approach is depicted in Fig. 2. One instance of the VQ-VAE is tasked to model the brains only, a second complementary instance models the full skulls *without* the brain. Applying both trained instances on an MRI dataset produces a paired dataset where each item is a tuple of the two latent integer volumes coming from each VQ-VAE encoder. In a second step, we employ a Paella-style MAE that conditions on the previously derived latents from the brain and models the distribution of the latents pertaining to the full skulls (see Fig. 3). Unseen MRI scans can then be de-identified by computing the latents associated to the brain and letting the MAE synthesize a realistic skull using the brain latent as conditioning variable (see Fig. 4).

³ As the skull is synthesized around the brain, $\gamma(\cdot)$ should at least contain a binary brain mask.




Fig. 3: **Latent Modeling Stage (MAE).** In the latent modeling stage, the skull and brain representations obtained in Fig 2 are provided to an iterative masked autoencoder (MAE), which repeatedly perturbs the skull latent and, using the brain latent as conditioning, reconstructs it.

Vector Quantization Stage. In order to reduce the memory requirements of high-resolution MR imagery, we leverage the VQ-VAE framework to compress (parts of) MRI scans into latent codes. As a preparatory step we use ROBEX [7] to compute the (binary) brain mask $B(X)$ of an MRI scan $X \in \mathcal{D}$. We need the brain representation for two reasons: it serves as a conditioning variable that informs the synthesis about the proportions of the skull, and we require it to later copy the original brain into the de-identified scan.

We then train a 3D VQ-VAE for each of the two representations $B(X) \odot X$ (*brain*) and $\bar{B}(X) \odot X$ (*skull*) *in isolation* where $\odot$ denotes the Hadamard product and $\bar{B}(X) \triangleq 1 - B(X)$. Both representations are complementary⁴ to each other in that $B(X) \odot X$ contains solely the brain (i.e., has non-zero brain intensities) without featuring any of the remaining parts (i.e., zero non-brain intensities) whereas the inverted properties are true for $\bar{B}(X) \odot X$.

Training both models yields two encoder/decoder pairs (e_1, d_1) and (e_2, d_2) . After training, we deploy the two encoders to translate the dataset of MRI scans $\mathcal{D}$ into a highly-compressed *paired* dataset of integer codes $\mathcal{D}_e = \{(e_1[B(X) \odot X]), e_2[\bar{B}(X) \odot X] \mid X \in \mathcal{D}\}$. A depiction of this and the applied preprocessing can be found in Fig 2.

Latent Modeling Stage. Having the compressed latent integer representations of $E_1 \triangleq e_1[B(X) \odot X]$ (*brain*) and $E_2 \triangleq e_2[\bar{B}(X) \odot X]$ (*skull*) now at our disposal, we describe a generative MAE $H_\Phi(E_2|E_1)$ that predicts the integer skull representation while conditioning on the integer brain representation.

Recall that both E_1 and E_2 are 3D volumes/grids of integers from the set $\{0, \dots, N_{CV} - 1\}^{s \times s \times s}$. In order to model the distribution of E_2 , the MAE methodology requires that one defines a time-dependent scheme to *perturb* the integers in the E_2 representation. We adopt *Paella*-style [15] perturbations, i.e., a value $v \in \{0, \dots, N_{CV} - 1\}$ is kept with probability α_t and resampled from the same set with a probability $1 - \alpha_t$, where the time $t \sim \mathcal{U}(0, 1)$ is sampled independently for each element in the batch.

Sampling Stage. After training the MAE, we execute the following sampling steps: (i) Compute the brain mask $B(X')$, (ii) compute the integer representation $E_1 = e_1[B(X') \odot X']$, (iii) use H_Φ to sample E_2 using E_1 as a conditioning

⁴ By adding both one recovers X




Fig. 4: **Test-time De-Identification.** We repeat the steps from Fig 2 to obtain the highly-compressed brain representation which is used as the *condition* in the inference stage of the masked autoencoder. Starting from a randomly-initialized $\hat{X}$ the network refines its estimate of how a skull around the given brain could look like. The final *de-identified* scan is obtained by *blending* the original scan with the last estimate $\hat{X}$ where the *binary* brain representation acts as a *mask*. This step ensures that the brain is preserved.

variable, and finally (iv), pass E_2 through the decoder $d_2(\cdot)$ to recover a skull that harmonizes well with the brain in X' . To ensure that the de-identified scan accurately represents the real brain and not a hallucinated version thereof, we implement a simple blending scheme in order to copy the brain from the original X via $Y = (1 - B(X)) \odot d_2(\hat{E}_2) + B(X) \odot X$ where $\hat{E}_2$ denotes the latent as estimated by letting the MAE condition on $E_1 \hat{=} e_1[B(X) \odot X]$ (i.e., the brain).

4 Experiments & Conclusion

We evaluate our MRI de-identification method’s effectiveness and ability to preserve the diagnostic value of the MRI retention by assessing both privacy protection and performance of downstream diagnostic tasks. This includes the segmentation tasks FIRST [12] and FASTSURFER [6]) to gauge the impact of de-identification on real-world applications, along with user and model-based studies to determine re-identification difficulty.

Setup. Our study utilizes the ADNI [21,22] (2,172 scans) and OASIS-3 [9] (2,556 scans) datasets, with a patient-wise 80-20 train-test split. For training our network, we leverage two NVIDIA RTX 4090 GPUs, each equipped with 24 GiB of GPU memory. Our implementation relies on the PyTorch framework [11]. We use the *inductor* backend for model compilation to optimize memory consumption.

Benchmarks. Our study compares against six de-identification methods: QUICK-SHEAR (removal), FACE MASK v1 (removal), DEFACE (removal), MRI WATER-SHED (removes non-brain tissue), FACE MASK v2 (removal), and CP-GAN (re-modeling) applied to 128^3 and 256^3 resolution images (CP-GAN is limited to 128^3).

<i>Defeat de-identification (Human accuracy)</i>				
	ADNI		OASIS-3	
	128 ³	256 ³	128 ³	256 ³
BLACK	20.17 ± 25.45	20.30 ± 26.33	18.37 ± 10.90	18.54 ± 11.08
BLURRED	45.05 ± 28.91	46.29 ± 29.42	41.63 ± 15.57	42.93 ± 18.06
ORIGINAL	55.33 ± 30.70	58.93 ± 29.24	61.86 ± 15.00	59.27 ± 16.94
MRI WATERSHED	19.03 ± 25.31	21.48 ± 27.25	22.56 ± 13.11	22.20 ± 14.41
DEFACE	43.53 ± 30.16	45.58 ± 29.53	38.14 ± 9.58	43.17 ± 16.04
QUICKSHEAR	38.51 ± 30.18	39.81 ± 30.43	40.70 ± 16.68	35.85 ± 13.78
FACE MASK v1	48.55 ± 30.49	50.24 ± 30.54	50.23 ± 17.39	52.68 ± 14.32
FACE MASK v2	38.05 ± 28.79	43.60 ± 29.90	33.02 ± 17.93	35.85 ± 15.00
CP-GAN	28.46 ± 27.95	✗	30.47 ± 14.30	✗
CP-MAE	23.14 ± 27.88	25.91 ± 27.65	22.56 ± 11.77	23.41 ± 15.27

<i>Defeat de-identification (Model accuracy)</i>				
	ADNI		OASIS-3	
	128 ³	256 ³	128 ³	256 ³
BLACK	19.66 ± 2.25	18.75 ± 2.08	19.86 ± 0.89	19.98 ± 1.19
BLURRED	87.46 ± 7.65	86.54 ± 8.61	97.05 ± 2.33	97.31 ± 1.67
ORIGINAL	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
MRI WATERSHED	44.75 ± 4.09	47.03 ± 5.06	67.76 ± 4.39	67.50 ± 7.39
DEFACE	99.27 ± 0.86	99.22 ± 1.51	99.78 ± 0.36	99.85 ± 0.20
QUICKSHEAR	98.79 ± 0.87	95.66 ± 1.98	99.81 ± 0.21	99.83 ± 0.20
FACE MASK v1	96.31 ± 2.71	98.75 ± 1.59	99.65 ± 0.50	99.71 ± 0.23
FACE MASK v2	94.42 ± 5.11	98.06 ± 1.20	99.78 ± 0.31	99.63 ± 0.29
CP-GAN	56.11 ± 5.05	✗	56.40 ± 2.89	✗
CP-MAE	39.91 ± 9.49	41.74 ± 6.91	48.82 ± 4.32	58.19 ± 4.04

Table 1: **De-Identification Quality.** We compare CP-MAE against traditional methods, CP-GAN and four control methods (above "- -" line) in terms of their de-identification capabilities w.r.t. a user-based (top) and a model-based (bottom) task. In both scenarios, given an MRI scan, the task is to identify the original scan among five de-identified alternatives (four of which belong to different subjects). We report the rate of correct guesses ($\pm$ s.d.), thus lower scores indicate better de-identification. "✗" indicates an unsupported resolution.

De-Identification Quality: User-based Study. In our Mechanical Turk study appearing in Table 1 we assessed the de-identification methods' resilience by having human subjects attempt to defeat the de-identification by matching 3D scans with de-identified versions from five choices, using side-profile views. Across 22,000 responses, CP-MAE excelled in de-identification, outperforming both traditional methods and CP-GAN by a substantial margin.

De-Identification Quality: Model-based Study. We also compared the privacy value of the de-identification models by using a neural network to attempt to defeat the de-identification in a manner similar to [5] but with a more powerful similarity-quantification network and side (45°) instead of frontal face views that significantly increase the level of detail. A Siamese network was trained to recognize patient-matched images, utilizing ResNet-18 for embeddings and




Fig. 5: **Downstream Tasks: Subcortical segmentation.** We analyze to which extent de-identification affects the quality of subcortical segmentation methods. The depicted values are the *class-averaged* Dice scores over 15 classes for FIRST and resp. 78 classes for FASTSURFER. To increase visual discernability we excluded MRI WATERSHED (avg. ≈ 0.21). Higher values are preferable.

computing distances with the Triplet Margin loss. Then, for each test fold, we present a patient’s original scan and its corresponding de-identified version (using method m) along with four incorrect options randomly chosen m -renderings from different patients. The network attempts to re-identify by selecting the option, y , minimizing the Euclidean distance to $\tilde{h}(x)$. The results in Table 1, show CP-MAE with the highest resistance to re-identification across the board. DEFACE and both FACE MASK models are more easily defeated as they mainly act on the face, enabling the network to leverage the information provided by the side view (e.g., the ears) to determine the right identity.

Effect of De-Identification on Medical Analyses. To ensure de-identification does not hinder medical analysis, we tested its impact using the subcortical brain segmentation algorithms FIRST [12], and FASTSURFER [6]. The original and de-identified scan segmentations are compared with the Dice coefficient. CP-MAE maintained analysis integrity, outperforming other de-identification methods on both tasks. In contrast, MRI WATERSHED significantly disrupted downstream analyses, underscoring the shortcomings of methods that simply retain the brain and discard everything else.

Conclusion. In this work, we introduce CP-MAE, a new technique to de-identify MRI scans by remodeling the skull and face with a generative model. Our approach combines dual instances of VQ-VAE with an MAE that operates in a discrete compressed latent space. This setup enables CP-MAE to stochastically generate high-quality MR scans with the original brain content, but supporting significantly altered external appearance. We enhance the 3D de-identification resolution capability of CP-GAN from 128^3 to 256^3 , achieving an eightfold voxel increase. Our experiments confirmed CP-MAE’s superior de-identification efficacy over competing methods. Furthermore, our analysis on standard brain analysis tasks reveals that CP-MAE minimally impacts diagnostic accuracy. Future explorations will extend to the CT imaging domain.

Acknowledgements. Data collection and sharing for this project was funded by ADNI (National Institutes of Health Grant U01 AG024904) and DOD ADNI (DOD award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer’s Therapeutic Research Institute at the University of Southern California (USC). ADNI data are disseminated by the Laboratory for Neuro Imaging at the USC.

Disclosure of Interests. The authors have nothing to disclose.

References

1. Bischoff-Grethe, A., et al.: Defacing method for magnetic resonance imaging data. *Journal of Neuroscience Methods* **350**, 109034 (2021)
2. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer (2022)
3. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
4. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems*. pp. 2672–2680 (2014)
5. der Goten, L.A.V., Hepp, T., Akata, Z., Smith, K.: Conditional de-identification of 3d magnetic resonance images (2021)
6. Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B.R., Reuter, M.: Fastsurfer - a fast and accurate deep learning based neuroimaging pipeline. *NeuroImage* **219**, 117012 – 117012 (2019)
7. Iglesias, J.E., et al.: Robust brain extraction across datasets and comparison with publicly available methods. *IEEE Transactions on Medical Imaging* **30**(9), 1617–1634 (2011)
8. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)

9. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A.G., Raichle, M.E., Cruchaga, C., Marcus, D.: Oasis-3: Longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. medRxiv (2019). <https://doi.org/10.1101/2019.12.13.19014902>, <https://www.medrxiv.org/content/early/2019/12/15/2019.12.13.19014902>
10. Milchenko, M., Marcus, D.: Obscuring surface anatomy in volumetric imaging data. *Neuroinformatics* **11**(1), 65–75 (jan 2013). <https://doi.org/10.1007/s12021-012-9160-3>, <http://www.ncbi.nlm.nih.gov/pubmed/22968671><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3538950>
11. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E.Z., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. *CoRR abs/1912.01703* (2019), <http://arxiv.org/abs/1912.01703>
12. Patenaude, B., Smith, S., Kennedy, D., Jenkinson, M.: A bayesian model of shape and appearance for subcortical brain segmentation. *NeuroImage* **56**, 907–22 (02 2011). <https://doi.org/10.1016/j.neuroimage.2011.02.046>
13. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding by generative pre-training. *OpenAI Blog* (2018)
14. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners. *OpenAI Blog* **1**(8) (2019)
15. Rampas, D., Pernias, P., Zhong, E., Aubreville, M.: Fast text-conditional discrete denoising on vector-quantized latent spaces (2022)
16. Schimke, N., Kuehler, M., Hale, J.: Preserving privacy in structural neuroimages. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol. 6818 LNCS, pp. 301–308. Springer, Berlin, Heidelberg (2011)
17. Segonne, F., Dale, A., Busa, E., Glessner, M., Salat, D., Hahn, H., Fischl, B.: A hybrid approach to the skull stripping problem in mri. *NeuroImage* **22**, 1060–75 (08 2004). <https://doi.org/10.1016/j.neuroimage.2004.03.032>
18. de Sitter, A., Visser, M., Brouwer, I., Cover, K., van Schijndel, R., Eijgelaar, R., Müller, D., Ropele, S., Kappos, L., Rovira, Á., et al.: Facing privacy in neuroimaging: removing facial features degrades performance of image analysis methods. *European Radiology* **30**, 1062–1074 (2020)
19. Smith, S.M., et al.: Fast robust automated brain extraction. *Human Brain Mapping* **17**(3), 143–155 (2002)
20. Sohl-Dickstein, J., Weiss, E.A., Maheswaranathan, N., Martin, A.: Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239* (2020)
21. Weiner, M.W., Veitch, D.P., Aisen, P.S., Beckett, L.A., Cairns, N.J., Green, R.C., Harvey, D., Jack, C.R., Jagust, W., Morris, J.C., Petersen, R.C., Salazar, J., Saykin, A.J., Shaw, L.M., Toga, A.W., Trojanowski, J.Q.: The Alzheimer’s Disease Neuroimaging Initiative 3: Continued innovation for clinical trial improvement (may 2017). <https://doi.org/10.1016/j.jalz.2016.10.006>, <http://www.ncbi.nlm.nih.gov/pubmed/27931796><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC5536850>
22. Wyman, B.T., Harvey, D.J., Crawford, K., Bernstein, M.A., Carmichael, O., Cole, P.E., Crane, P.K., Decarli, C., Fox, N.C., Gunter, J.L., Hill, D., Killiany, R.J., Pachai, C., Schwarz, A.J., Schuff, N., Senjem, M.L., Suhy, J., Thompson, P.M., Weiner, M., Jack, C.R.: Standardization of analysis sets for reporting results from ADNI MRI data (may 2013). <https://doi.org/10.1016/j.jalz.2012.06.004>