Face Recognition and De-Identification of Research Brain Images with mri_reface

Christopher G. Schwarz
schwarz.christopher@mayo.edu
2023/05/23

The risk

- Public sharing of research data, including imaging, is being widely promoted and sometimes required.
- Typical de-identification of brain MRI removes only text “meta-data” stored in the image file.
  - Imagery of faces remains intact.
  - Is this de-identified *enough?*
- HIPAA standard for de-identification require removing:
  - “Full-face photographs and any comparable images”[1]

Question: Does brain MRI count as a comparable image?

MAYO CLINIC

[1] HIPAA §164.514(a)

Possible motives for re-identification

- As innocent as:
  - My dad is in this study. Can I find out his amyloid status?
  - I’m in this trial. Am I receiving the intervention?
- As invasive as:
  - Employers choosing who to hire/fire
  - Corporations mining medical records to sell targeted advertising
- As malicious as:
  - Agents seeking medical information to discredit or blackmail political or corporate foes

MAYO CLINIC

Typical Face Recognition Problem

- Step 1: Provide examples of a “training set” of photographs of faces to be recognized

Bill Clinton

George Bush

Dick Cheney

Step 2: Given a “test” photograph of an unknown individual, match it to the correct face to identify them

Unknown person

George Bush

MAYO CLINIC

MRI Face Recognition Hypothesis

- Step 1: Provide examples of a “training set” of MRI-based reconstructions of faces to be recognized

Participant #41040

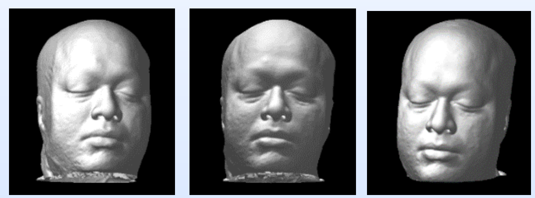

Participant #52352

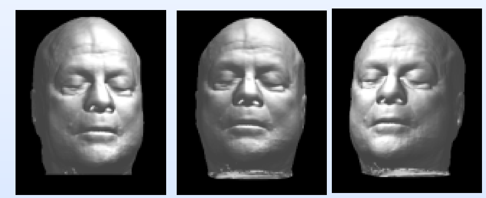

Participant #22072

Step 2: Given a “test” photograph of a named individual, match it to the correct MRI to identify their study data

John Smith

AKA

Participant #52352

Simulates a “population-to-sample attack”

Data Set

- **84** Mayo Clinic healthy volunteers, age 34-89
- Each had a previous head MRI within 3 months as part of existing enrollment in Mayo Clinic MCSA or ADRC
  - Siemens 3D FLAIR identical to ADNI3 protocols
- Captured 5 photographs of each volunteer with a standard iPad
  - These were the “test” images to be recognized.

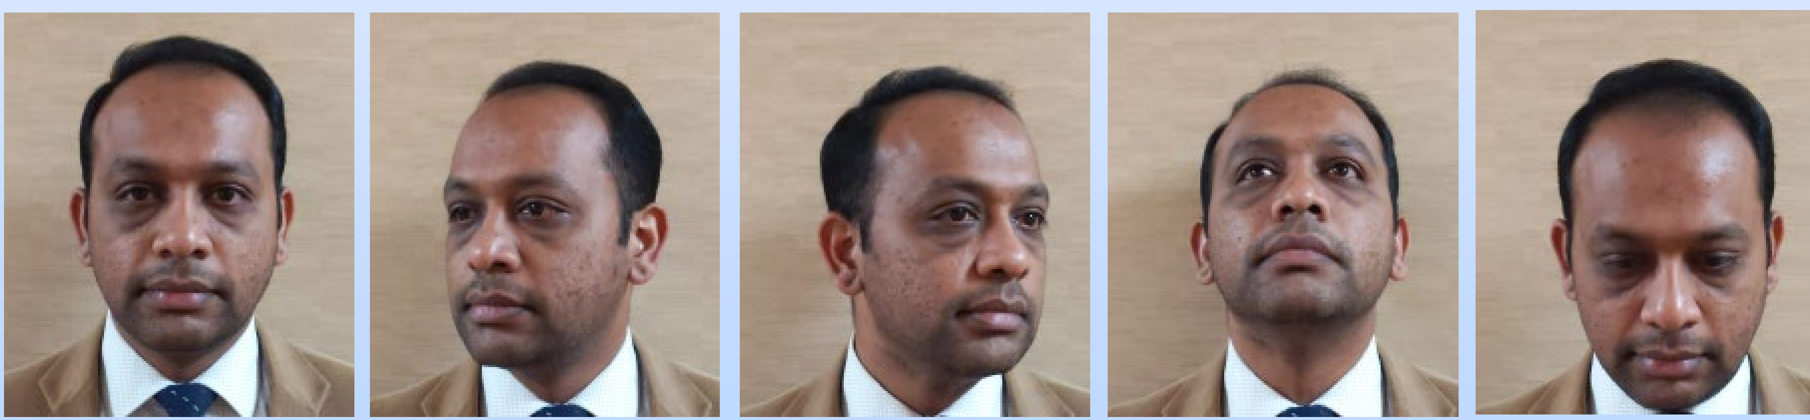

MAYO CLINIC

MRI-based face reconstructions

- Generated 81 2D photograph-like images of each MRI using Surf Ice[1], under varying simulated lighting positions and views.
- These were the “training” set of faces to be recognized.

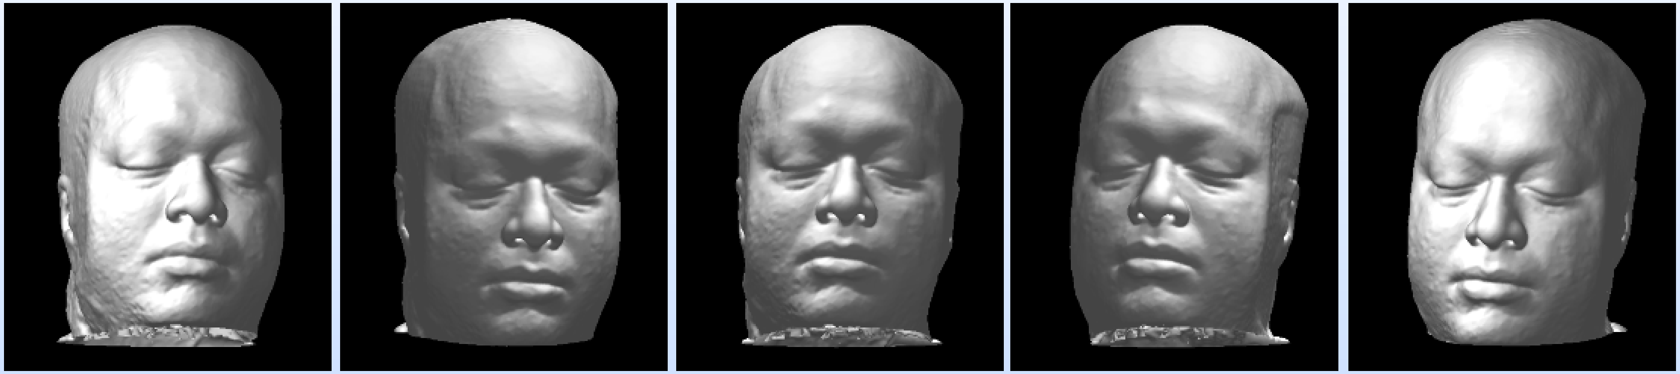

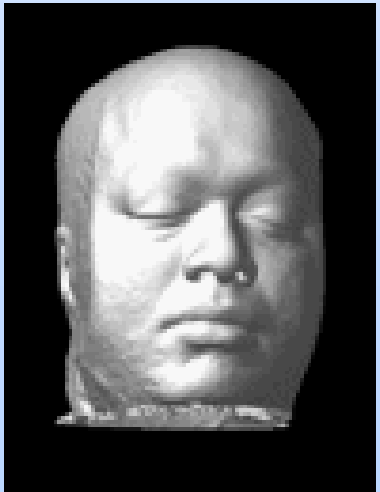

[1] https://www.nitrc.org/projects/surfice/

MAYO CLINIC

Testing: Can AI match photos to the correct MRI?

- We tested using software available easily and freely to general public (within a private, secure Mayo Clinic-owned cloud instance to protect the data)

Microsoft Azure Contact Sales: 1-800-867-1389 Search My account Portal Christopher

Overview Solutions Products Documentation Pricing Training Marketplace Partners Support Blog More Free account

Face identification

Face API enables you to search, identify, and match faces in your private repository of up to 1 million people.

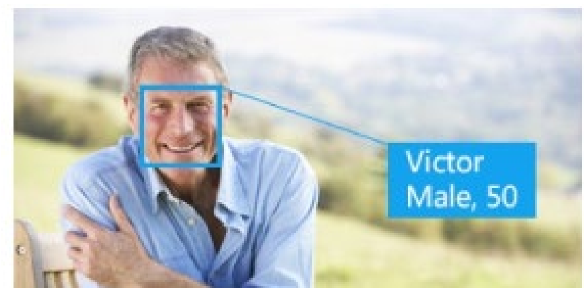

Pricing details

TIER	FEATURES	PRICE
Face API – free	Up to 20 transactions per minute	30,000 transactions free per month

Results

- For 70/84 participants (83%), their correct MRI (FLAIR) was chosen as the software's #1 match for their five photos
  - Matching by random chance would expect only one correct match. ($p<0.001$)
  - This was our study published in NEJM Oct 2019.
- Since then (2022), we have replicated with 182 people, and also tested PET/CT:

MAYO CLINIC

Schwarz et al. NEJM 2019

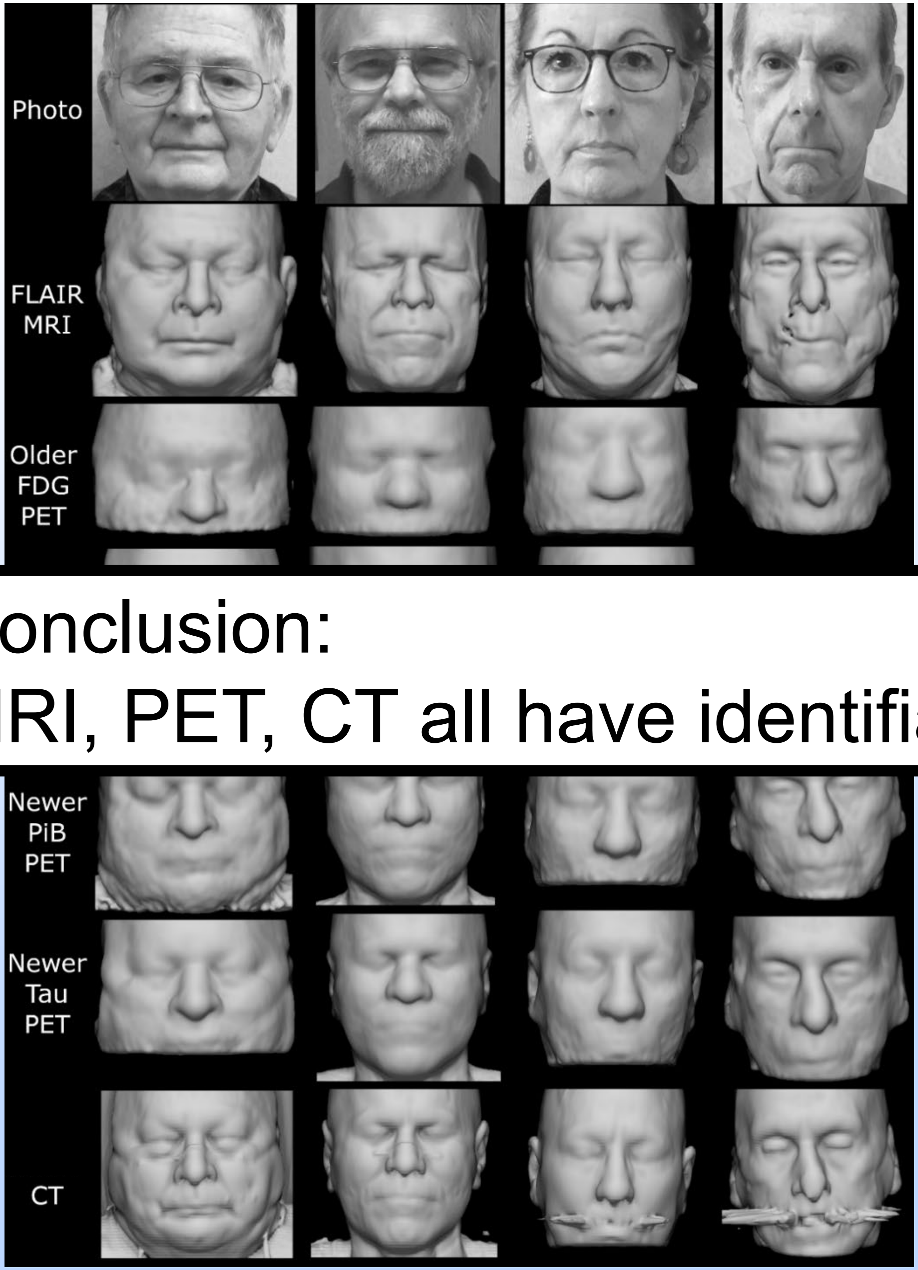

Match rates:

FLAIR: 178/182 (98%)
T1: 176/182 (97%)

Older FDG: 54/129 (41%)
Older PIB: 54/167 (32%)
Older FTP: 59/167 (35%)

Conclusion:
MRI, PET, CT all have identifiable faces!

Newer PIB: 17/20

Newer FTP: 18/19

CT (from older PET/CT):
131/167 (78%)

Adapted from Schwarz et al., NeuroImage 2022

Match rates across MRI sequences

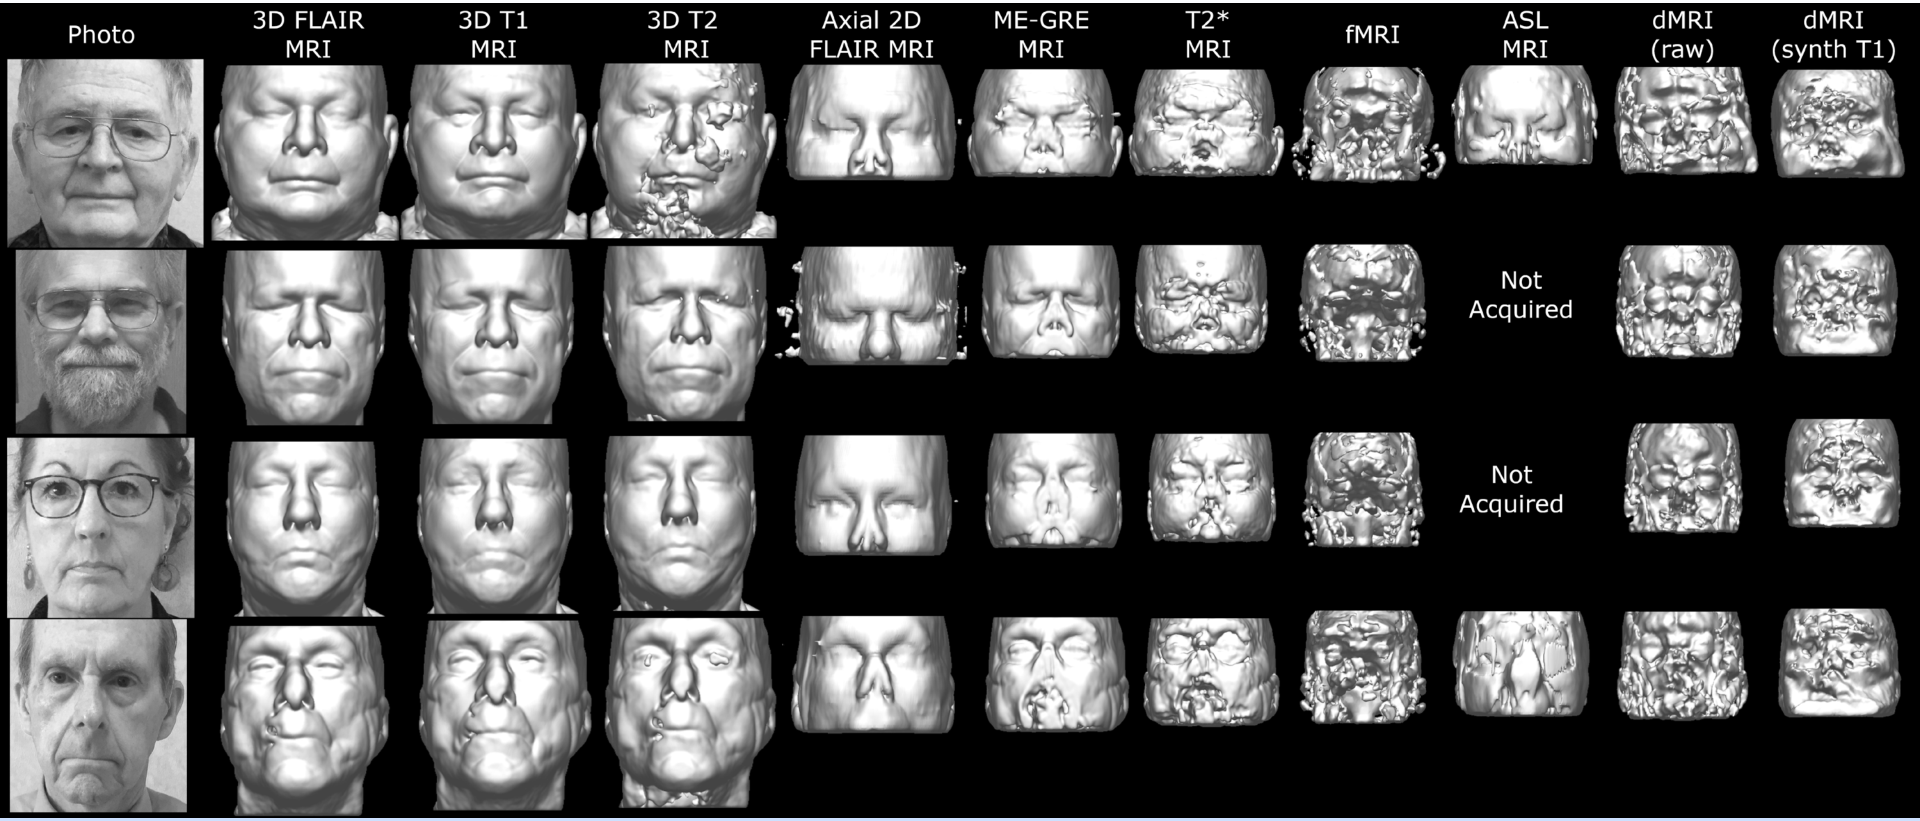

98% 97% 96% 44% 45% 10% 0% 8% 2% 1%

- Must de-face structural sequences
  - Larger risk for 3D than 2D
- EPI sequences (dMRI, fMRI, ASL) have minimal risk

From manuscript under revision

Related Literature

- Only one prior study using MRI: 40% of human visual raters could match MRI to photos with success rates above chance.[1]

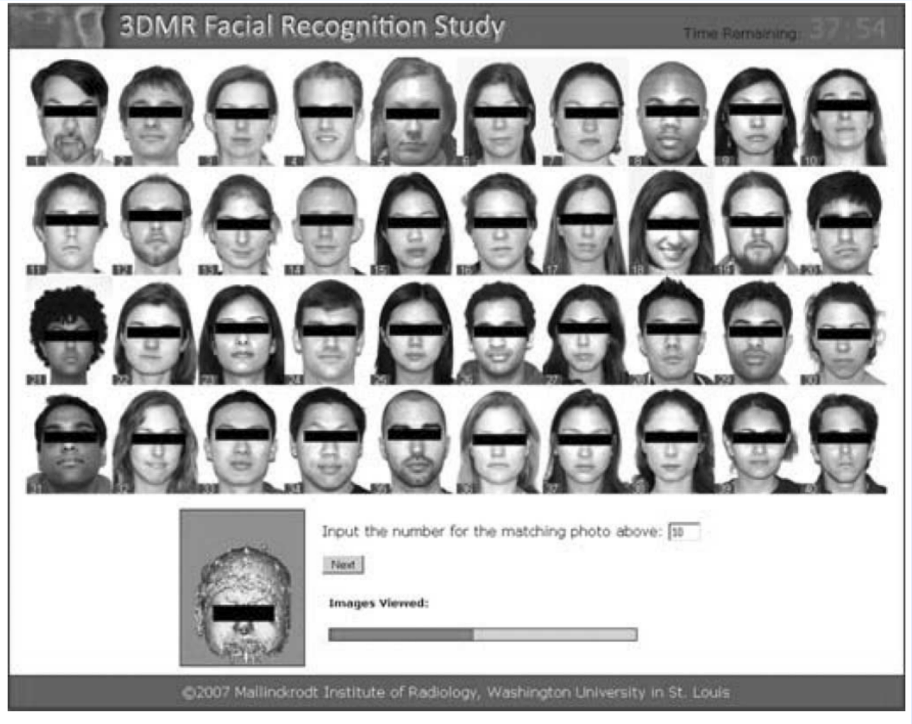

[1] Prior et al. IEEE Trans Inf Technol Biomed 2009

MAYO CLINIC

Related Literature

- Only one prior study using MRI: 40% of human visual raters could match MRI to photos with success rates above chance.[1]
- Only one prior study using automated face recognition: CT imaging match rate of 27.5% using Google Picasa (developed in 2008).[2]

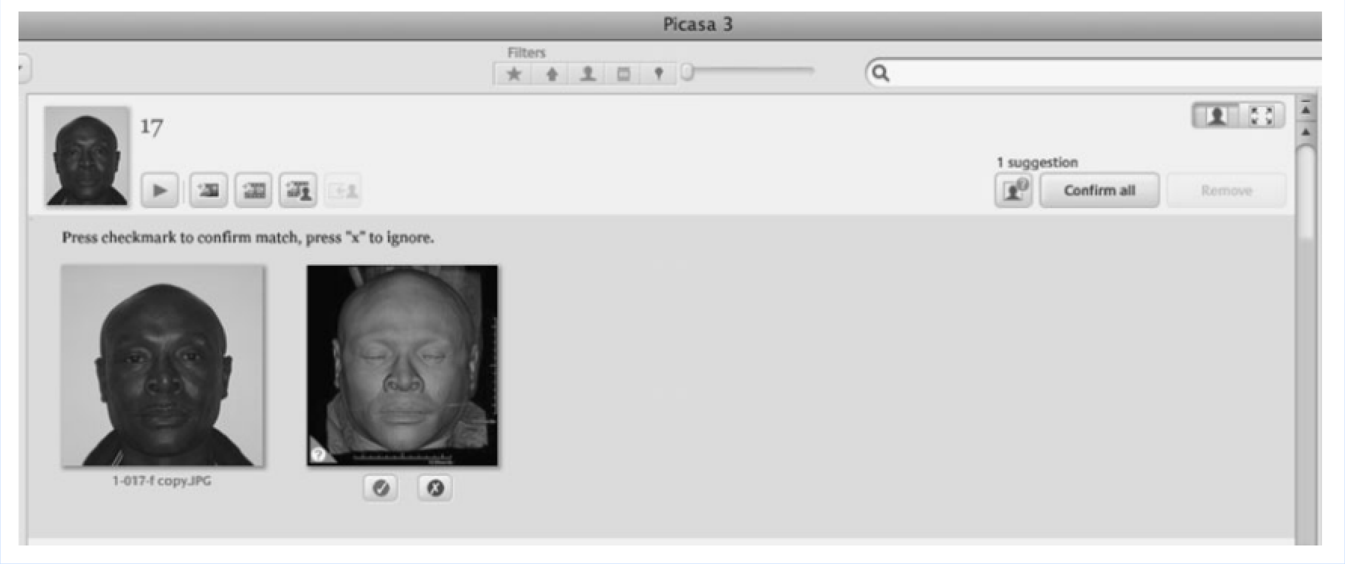

[1] Prior et al. IEEE Trans Inf Technol Biomed 2009
[2] Mazura et al. J Digit Imaging 2012

Related Literature

- Only one prior study using MRI: 40% of human visual raters could match MRI to photos with success rates above chance.[1]
- Only one prior study using automated face recognition: CT imaging match rate of 27.5% using Google Picasa (developed in 2008).[2]
- Our >90% match rate likely reflects recent advances in face recognition.
  - Deep learning-based methods have improved face recognition approximately 20x in the past 5 years.[3]

[1] Prior et al. IEEE Trans Inf Technol Biomed 2009
[2] Mazura et al. J Digit Imaging 2012
[3] Grother et al. 2018.

MAYO CLINIC

Part 2: Face De-identification (De-facing)

- Software to remove identifiable facial features in MRI ("de-facing") has been available since the mid-2000s, but typically has not been applied in U.S. aging studies.
  - Was believed that re-identification from face recognition would have a low success rate or wouldn't be attempted
  - Was concern that removing the face would hinder analyses
  - Existing software had minimal validation: it was assumed that removing at least the lower face, without removing any brain imagery, was good enough
- De-facing has generally been more prevalent in European studies, and in studies with younger populations.
- Cancer imaging datasets often have additional reasons not to de-face, as the pathology of interest may be outside the brain or face-adjacent.

MAYO CLINIC

Part 2: Face De-identification (De-facing)

- We can use an atlas to identify face voxels, but how exactly do we alter them to prevent recognition?

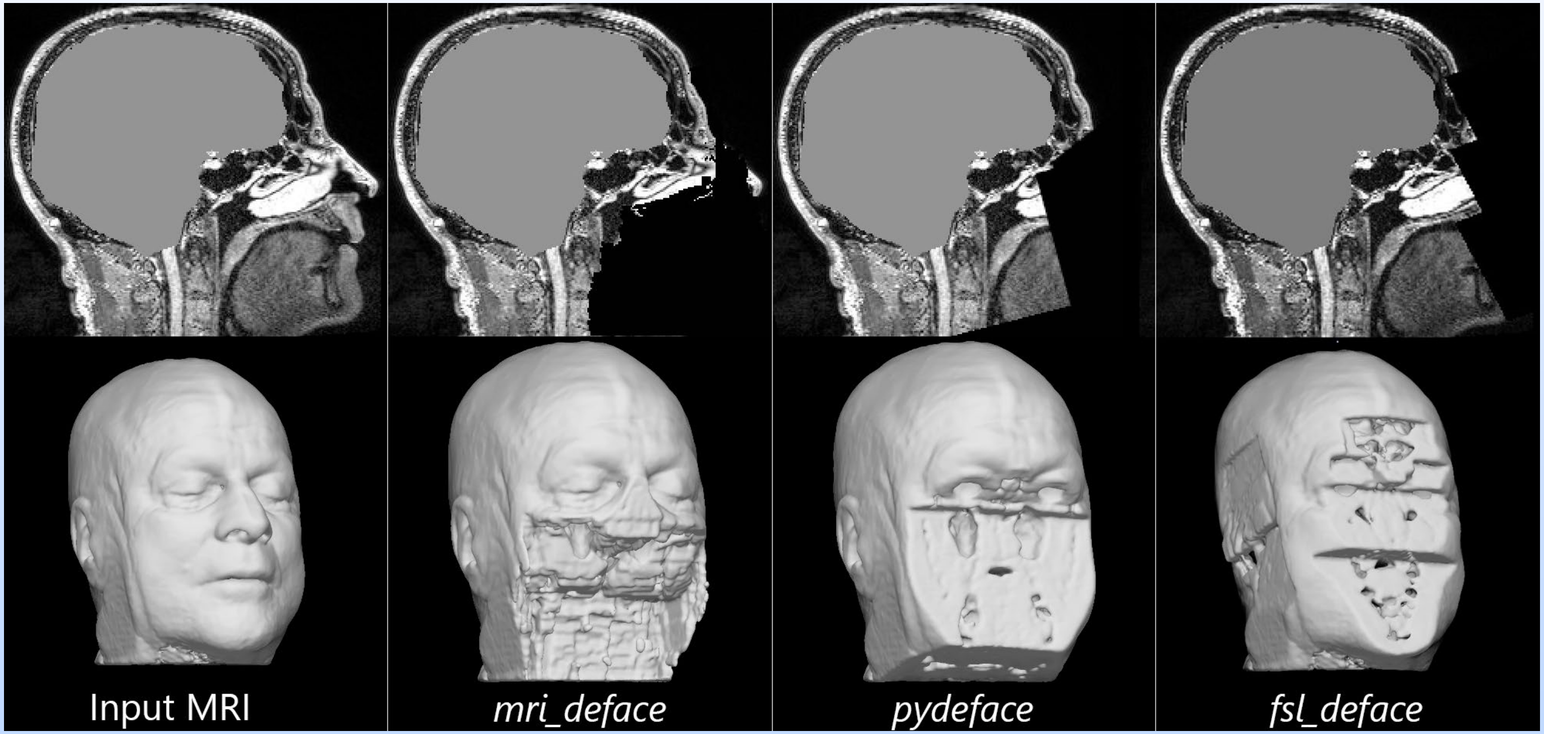

Schwarz et al. NeuroImage 2021

Part 2: Face De-identification (De-facing)

- mask_face (used in some HPC data) blurs only the outer face contour
- It can be automatically un-blurred to recover the original face[1].

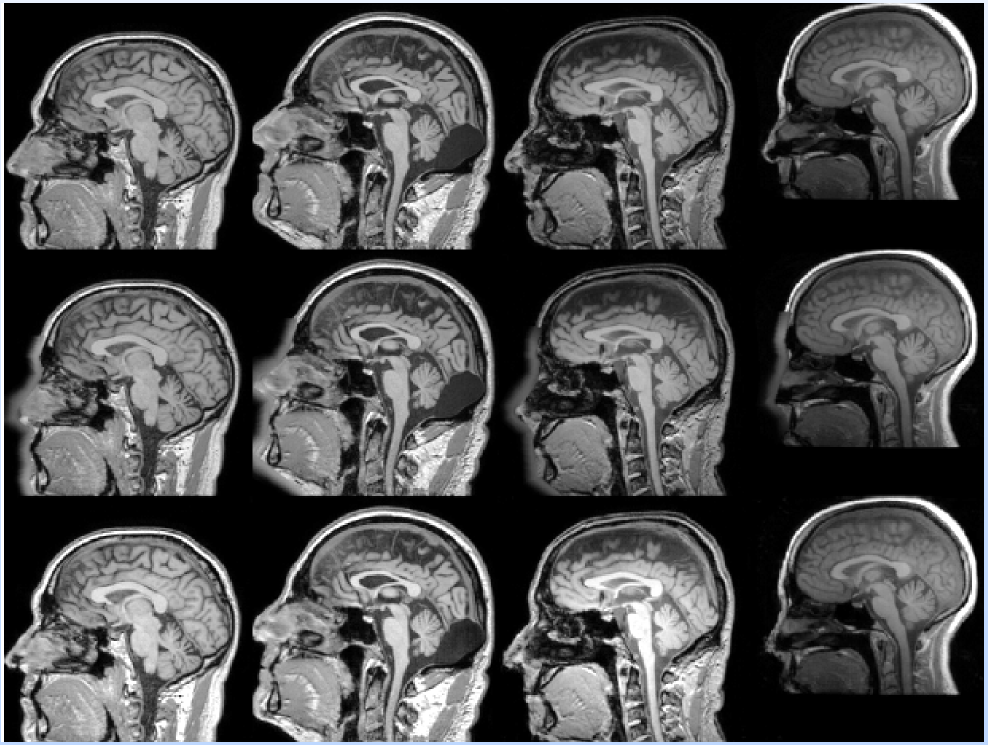

Ground Truth

De-identified

Reconstructed

MAYO CLINIC

Abramian et. al. Proc Int Symp Biomed Imaging 2019

mri_reface

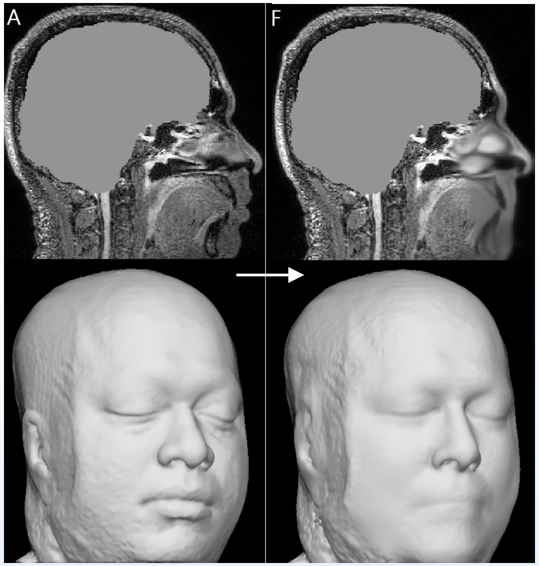

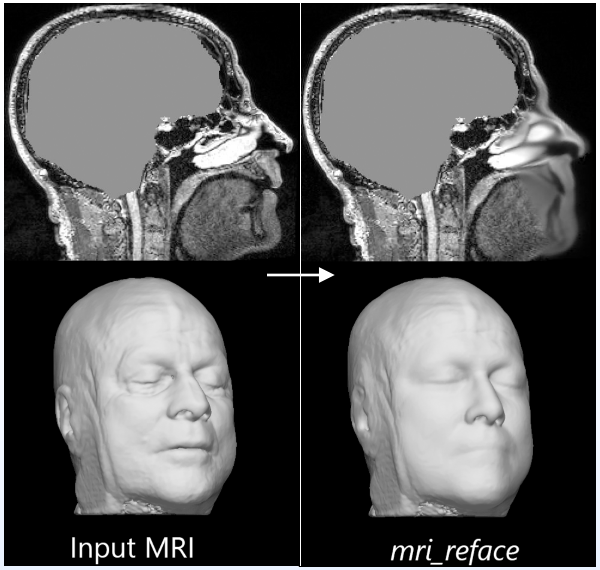

- Replace the face with an average face, rather than remove it
  - Produces a natural-looking image to reduce effects on pipelines downstream
- Nonlinear registration: can more precisely find face voxels.
- Also replaces ears, teeth, and aliased/wrapped face parts in front or behind the head
- Supports T1, T2, T2*, FLAIR, FDG PET, Amyloid PET, Tau PET, CT
- New: replacement face matches noise properties of input image
- Free on NITRC: **https://www.nitrc.org/projects/mri_reface**

MAYO CLINIC

Schwarz et al. NeuroImage 2021

Face recognition performance after de-facing

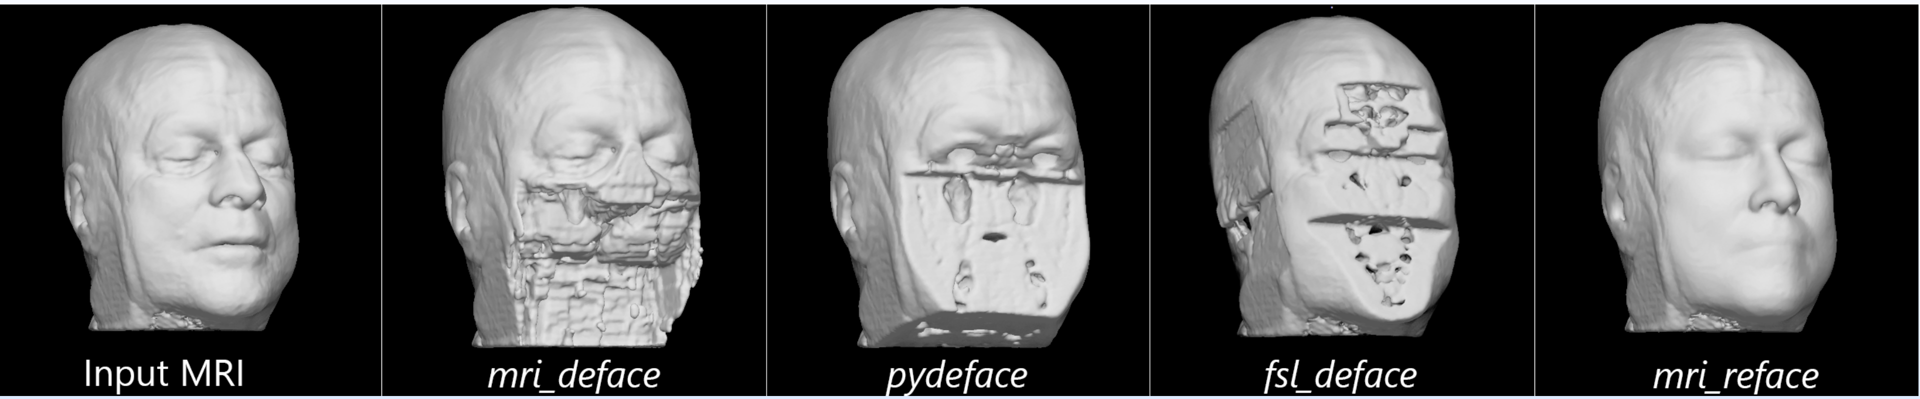

153/157	52/157	59/157	44/157	11/157
97%	33%	38%	28%	7%

- 7% still exceeds chance (1%), but far better than 97%
- We are continuing to improve our techniques

MAYO CLINIC

Updated from Schwarz et al. NeuroImage 2021

Effects of de-facing on GM volumes

- We ran each de-facer on 300 ADNI T1's
  - 100 from each vendor: 50 CU, 50 clinical AD
- Ran SPM12, FSL (UK-Biobank pipeline), and FreeSurfer 6.0 on each original and de-faced image.
- Compared regional GM volumes before/after

MAYO CLINIC

Schwarz et al. NeuroImage 2021

Effects of de-facing on GM volumes

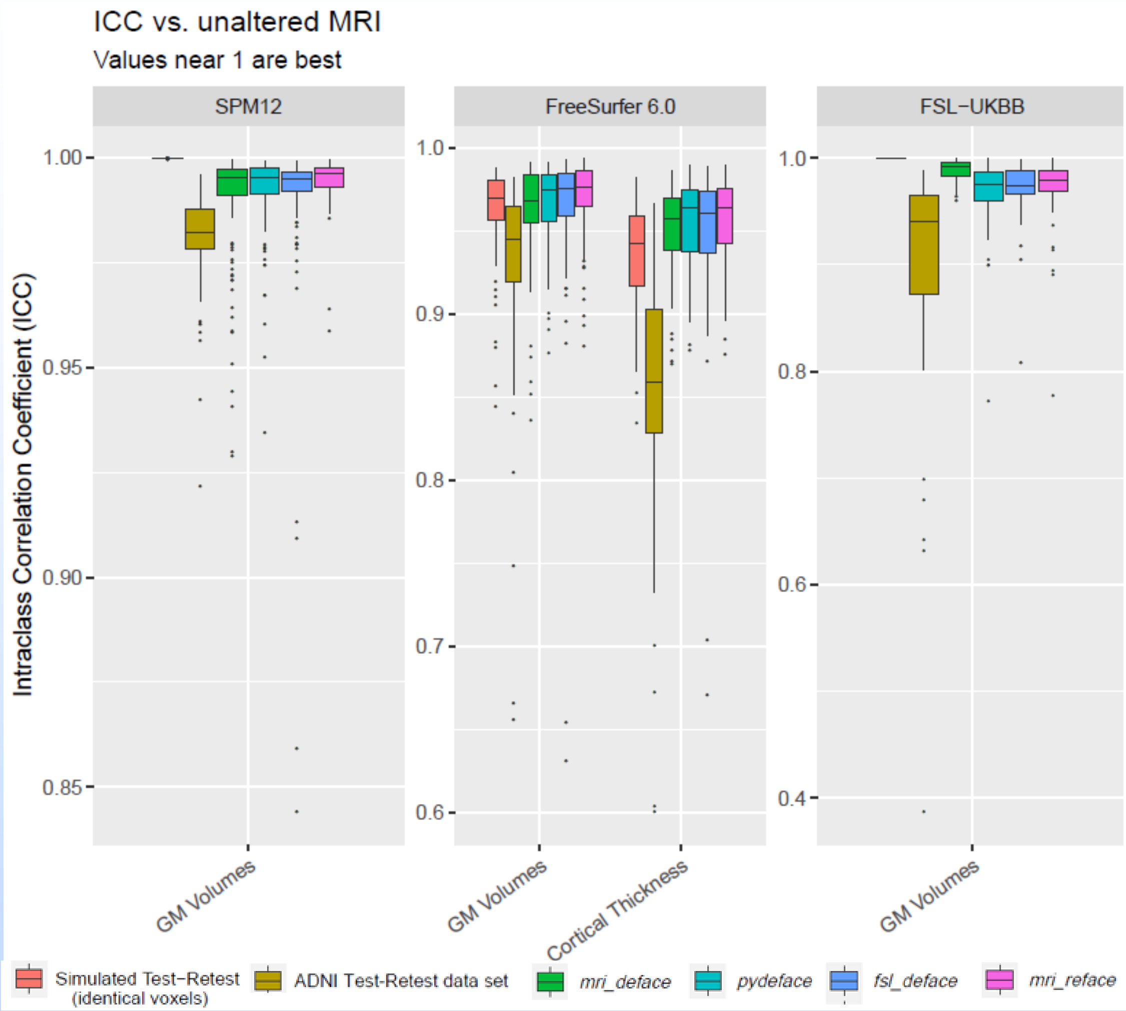

- Smaller than scan-rescan
- But larger than you'd think, given that all brain voxels are completely unchanged.

MAYO CLINIC

Schwarz et al. NeuroImage 2021

Effects of de-facing on GM volumes

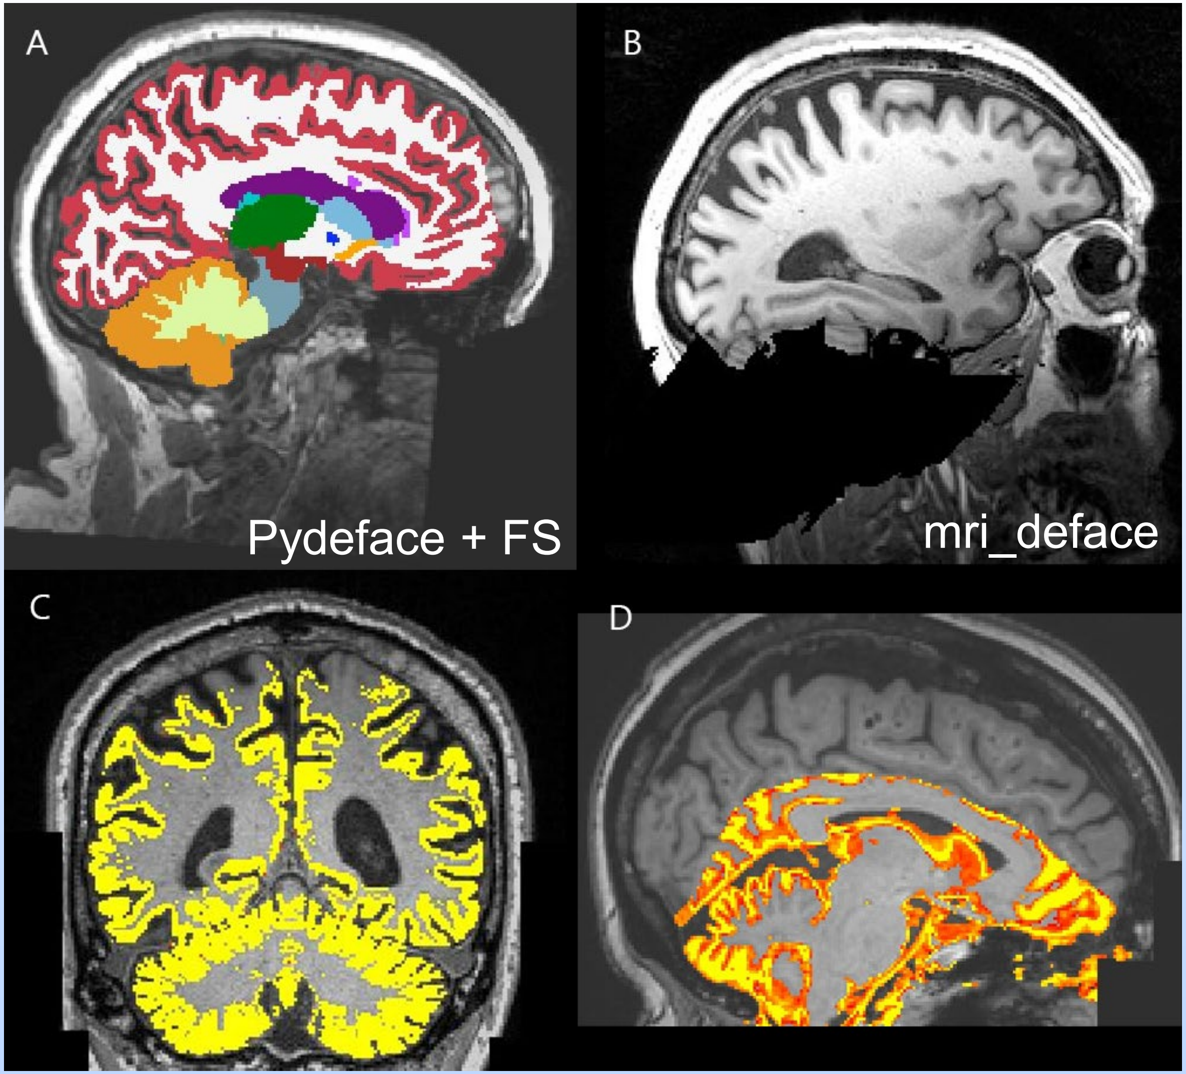

fsl_deface + SPM12

fsl_deface + FSL

Schwarz et al. NeuroImage 2021

A third-party comparison:

Effects of Defacing Whole Head MRI on Neuroanalysis

Chenyu Gao[a,*], Linghao Jin[b,*], Jerry L. Prince[a,b,c], and Aaron Carass[c]

[a]Department of Biomedical Engineering, The Johns Hopkins University, Baltimore, MD 21218
[b]Department of Computer Science, The Johns Hopkins University, Baltimore, MD 21218
[c]Department of Electrical and Computer Engineering,
The Johns Hopkins University, Baltimore, MD 21218

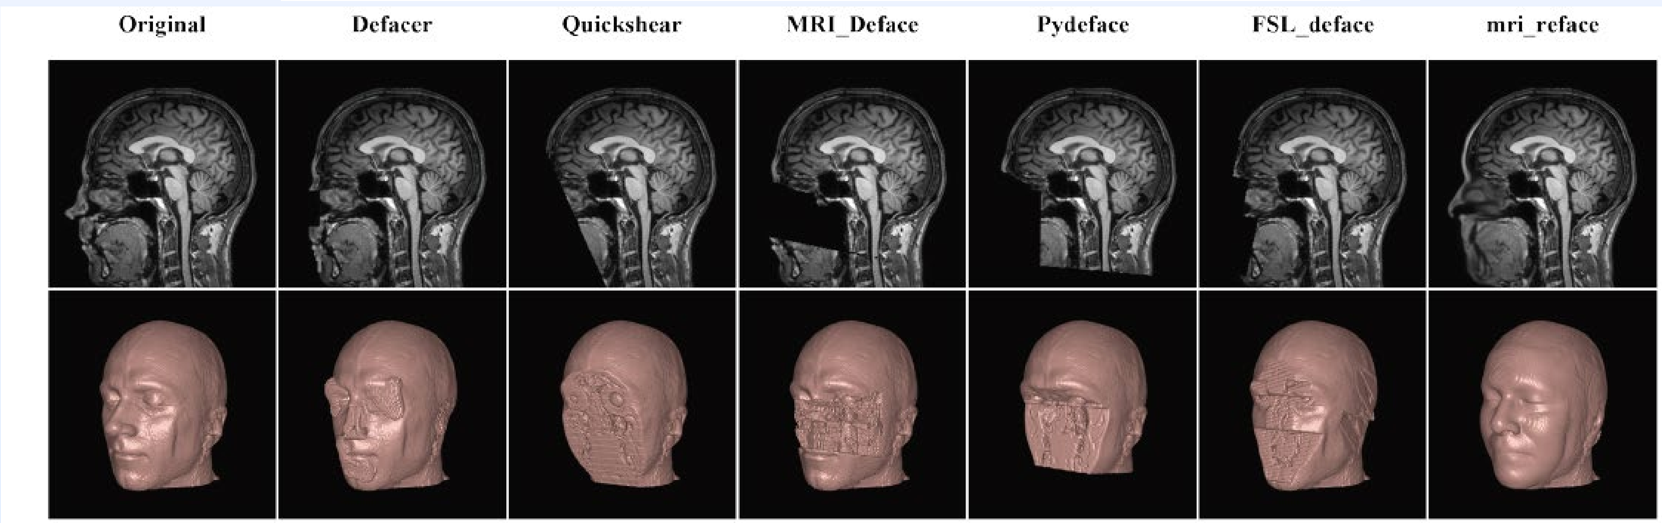

C. Gao et al., Proc SPIE 2022

A third-party comparison

Table 1. The results of the quality check. See section 3.2 for the definition of Success, Failure I, and Failure II.

Algorithm	Total	Success	Failure I	Failure II
Defacer	179	163	16	0
FSL_deface	179	142	0	37
MRI_Deface	179	136	7	36
mri_reface	179	179	0	0
Pydeface	179	179	0	0
Quickshear	179	78	0	101

- Most had high failure rates, >20% kept some face (failure 1) , or removed some brain (failure 2)
- *mri_reface* was one of only two methods that never failed in either way

MAYO CLINIC

C. Gao et al., Proc SPIE 2022

A third-party comparison

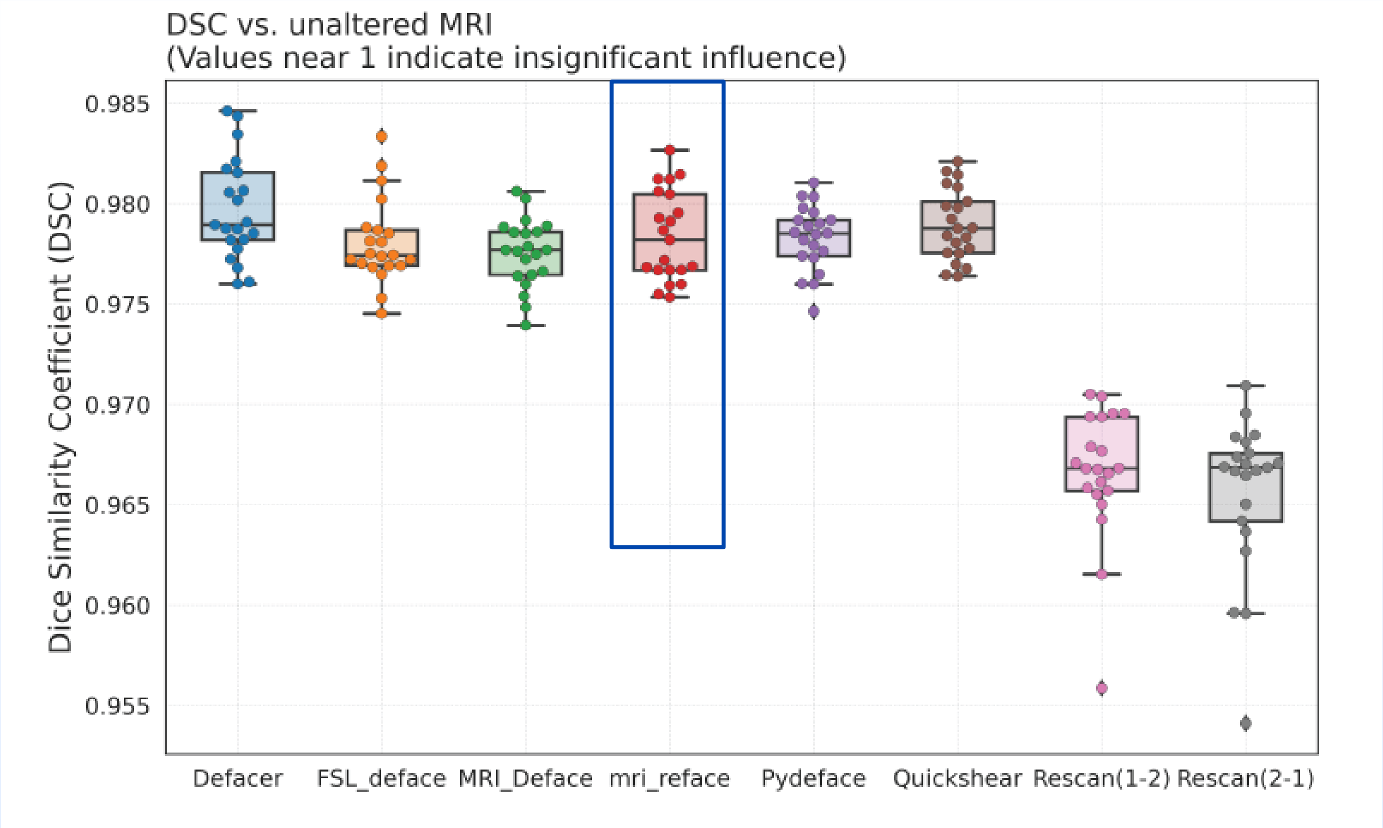

- All >> scan-rescan. Agrees with our findings.
- mri_reface roughly tied with 3 other de-facers:
  - On the previous slide, **Defacer** retained face in 9%, and **Quickshear** removed brain in 57%

MAYO CLINIC

C. Gao et al., Proc SPIE 2022

A third-party comparison

- The other winner so far was *pydeface*, but it performed the worst in most of their last set of comparisons.

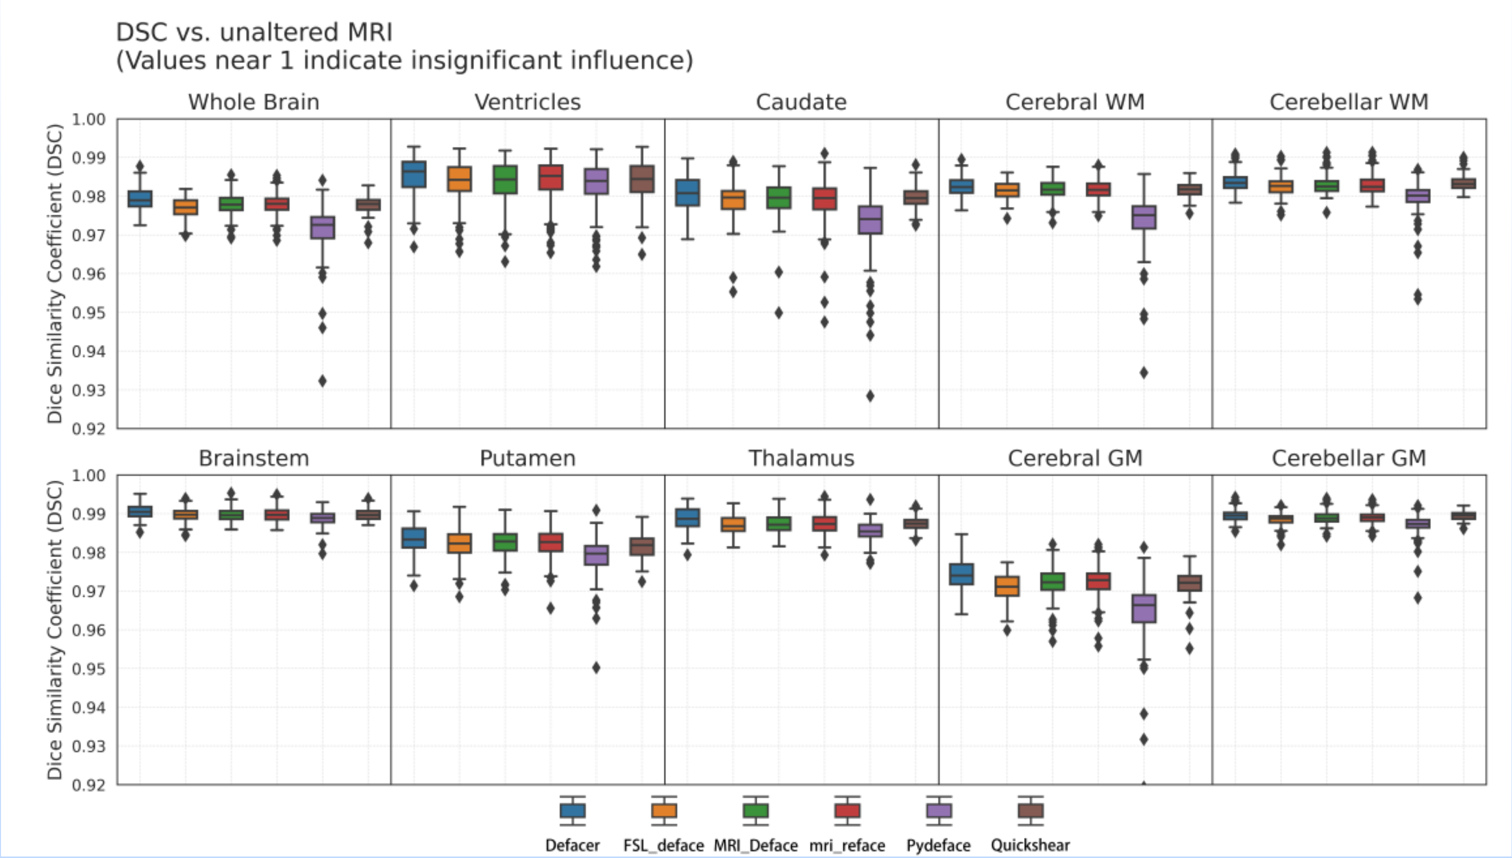

- In total, mri_reface had some of the smallest effects on measurements, and was one of only two to never fail.

MAYO CLINIC

C. Gao et al., Proc SPIE 2022

A second third-party comparison

Rubbert *et al. Insights into Imaging* (2022) 13:54
https://doi.org/10.1186/s13244-022-01195-7

Insights into Imaging

ORIGINAL ARTICLE **Open Access**

Impact of defacing on automated brain atrophy estimation

Christian Rubbert[1*], Luisa Wolf[1], Bernd Turowski[1], Dennis M. Hedderich[2], Christian Gaser[3], Robert Dahnke[3,4,5], Julian Caspers[1] and for the Alzheimer's Disease Neuroimaging Initiative

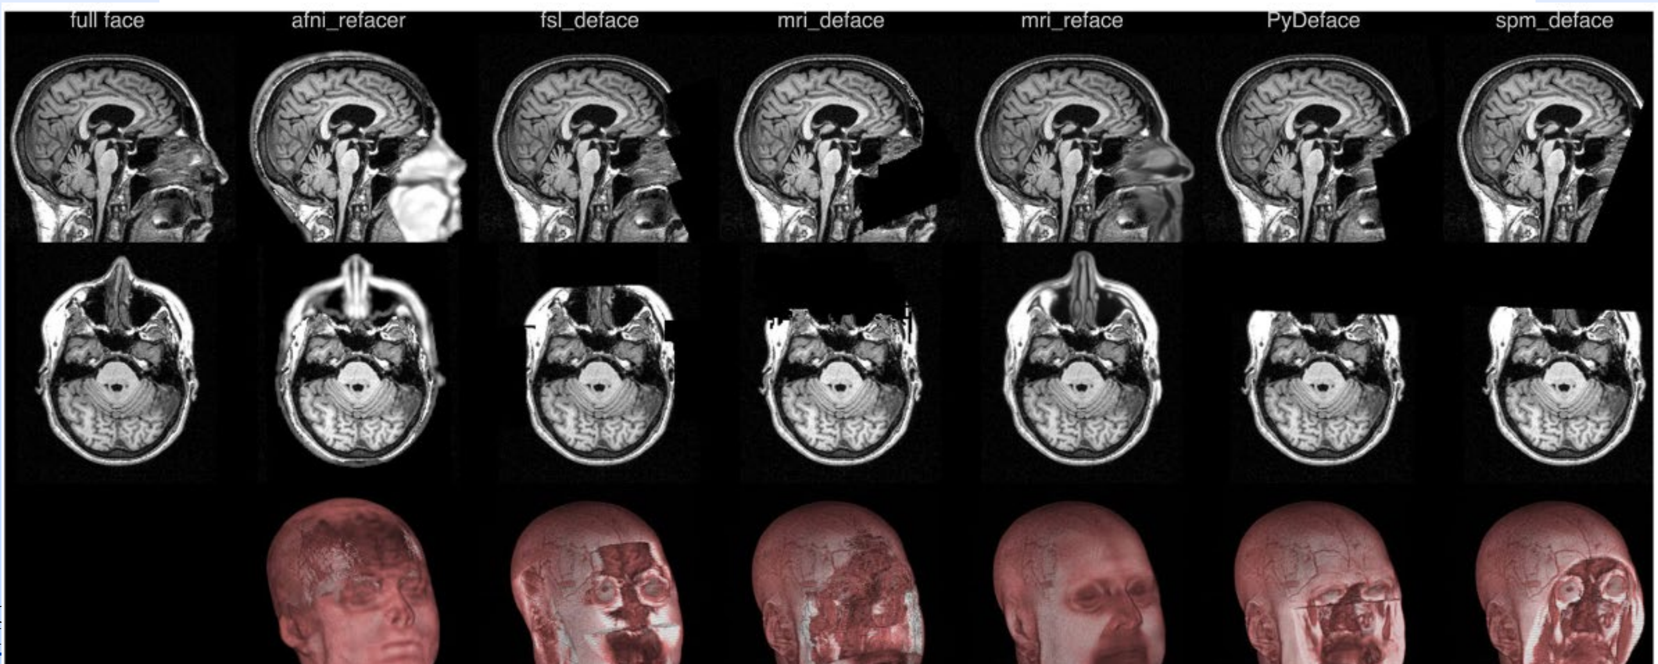

A second third-party comparison

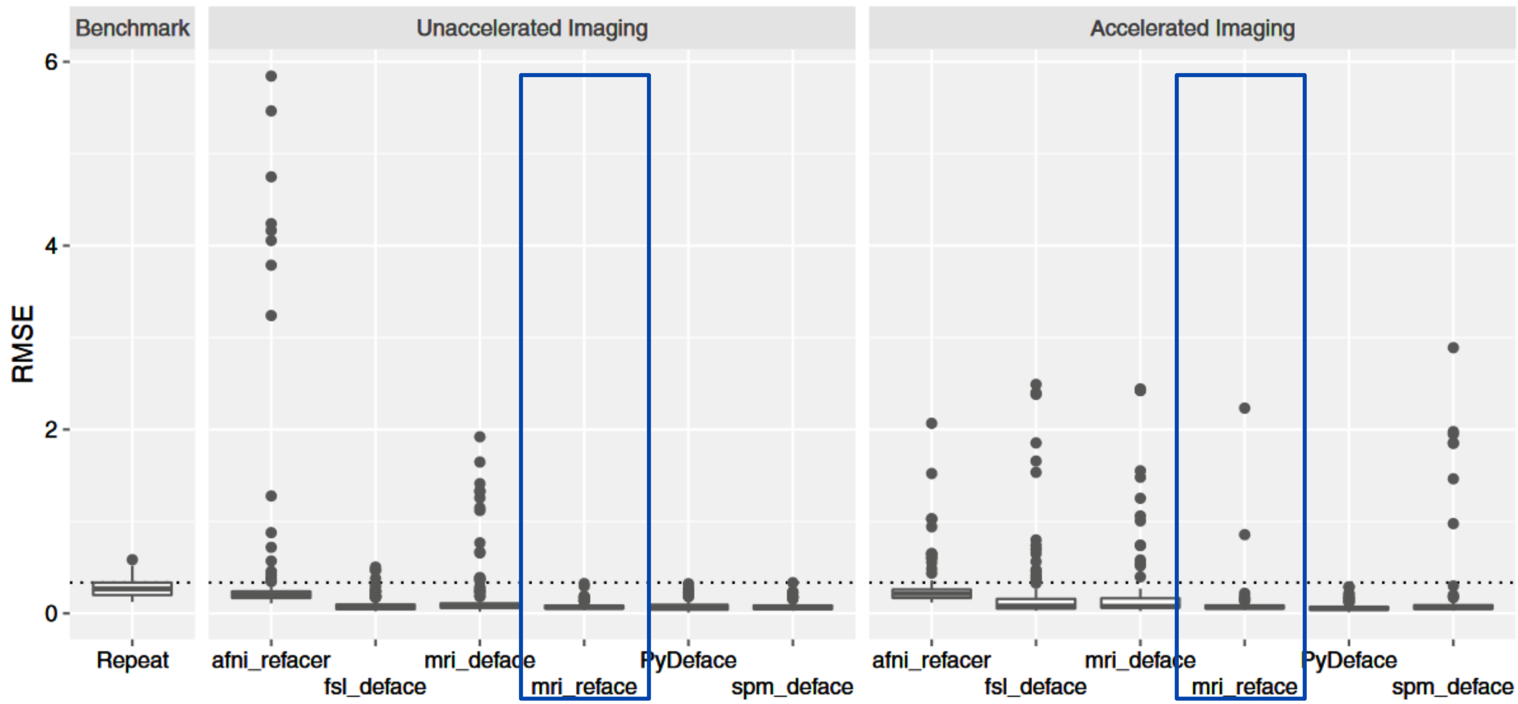

Fig. 2 Box-and-whisker plots of the root-mean-square error (RMSE) values after veganbagel processing of Alzheimer's disease patients. Left column: The benchmark result obtained by calculating the RMSE between gray matter z-scores of full face unaccelerated 3D T1 imaging and the respective same-session unaccelerated repeat 3D T1 imaging. Center and right: RMSE values obtained by comparing z-scores of defaced unaccelerated (center) or accelerated (right) 3D T1 imaging series with the respective full face 3D T1 imaging. The dotted black line denotes the 75th percentile of the RMSE values of the benchmark

- *mri_reface* had smaller effects than everything except *PyDeface,* which we've shown can still be recognized for 38%

MAYO CLINIC

Why does de-facing face affect brain measurements?

- Affine registration to template is affected by face, especially nose.
  - Affects the entire image
- Eyebrow ridge is most important to face recognition. It's also within a few MM of frontal lobe.
  - Largest effects in orbitofrontal, frontal pole

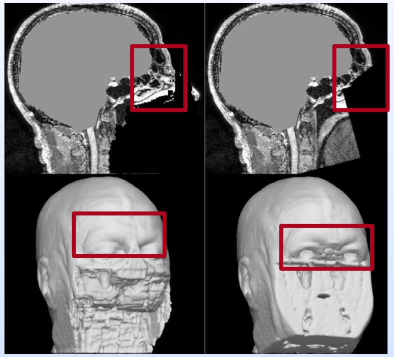

- Bayesian tissue class segmentations compare everything to everything else
  - Change in the not-GM intensity distribution affects probability of being GM
  - Largest effects in regions that are most difficult to segment normally: deep gray and sensorimotor

MAYO CLINIC

Schwarz et al. NeuroImage 2021

What about PET/CT?

- The latest version of *mri_reface* includes support for CT and PET (templates for FDG, PIB, AV45, FTP)
- Reduces PET recognition rates from 32-41% to 0-3.5%
- Reduces CT recognition rates from 78% to 5%
- Reduces MRI recognition rates from 97-98% to 8%
- Effects of de-facing on PET SUVR are measurable but negligible, like effects on T1 GM volume/thickness.

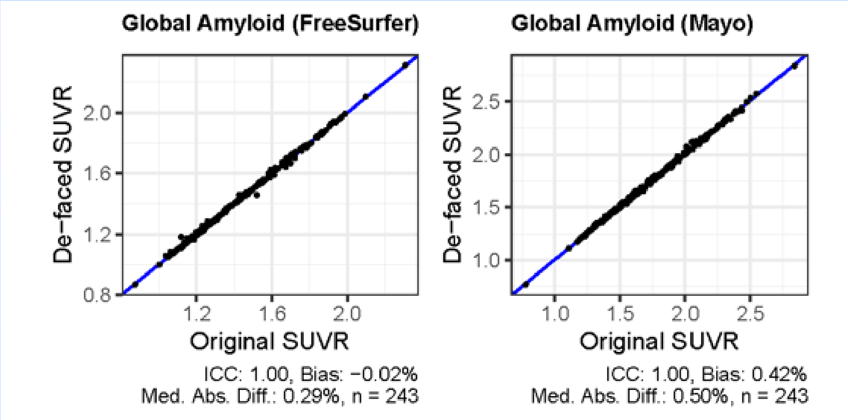

MAYO CLINIC

Adapted from Schwarz et al., NeuroImage 2022

Does de-facing affect correlations with clinical variables?

- We computed biomarkers of AD from MRI and PET with and without *mri_reface*, then compared the strengths and qualitative findings of their correlations with age and cognition.
- Findings had very high agreement between de-faced and unmodified images.
  - Voxel-wise comparisons were not significantly different
  - Among region-wise comparisons, only 3/55 correlations were significantly different, and these were not significant after correction for multiple comparisons.
  - Directions of not-significant differences were mixed: no predominance of weaker or stronger with de-faced images.
- For AD imaging research, de-facing with *mri_reface* had no effect on analyses.

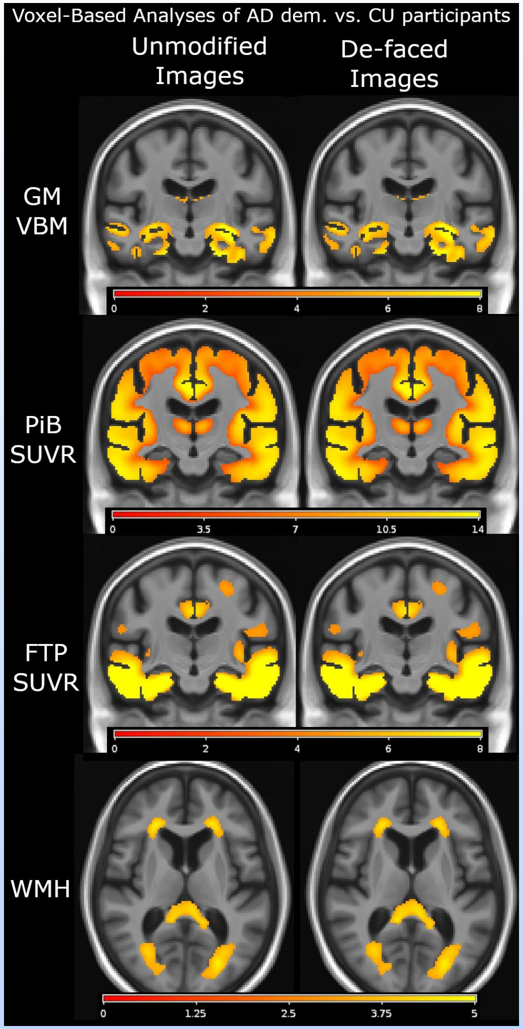

Manuscript under review

MAYO CLINIC

Limitations of de-facing

- *mri_reface* has minimal effects on biomarker measurements inside the brain, *but*:
  - Removing/replacing areas outside the brain can limit some secondary re-uses of the data
  - Some applications rely on stereotactic markers on the head, use the shape of the face for placement of leads, or use the complete volume of tissue in the head for dosage calibration.
  - Many cancer imaging datasets cannot readily apply de-facing because it would remove the pathology of interest.
- Future work will measure the effects of partial de-facing, by retaining some course face-parts or retaining a radius around marked pathology, on re-identification ability.
  - The eyebrow ridge is the most critical, but it may be possible to retain more nose/mouth with only mildly increased risk.

MAYO CLINIC

Conclusions

- Face recognition match rates are high enough (up to 98% in population-to-sample) to warrant some caution before unrestricted sharing of face-intact images.
- *mri_reface* can greatly reduce these match rates for MRI, PET, and CT (down to ≤8%) with negligible effects on brain measurements (<< scan-rescan differences).
- mri_reface is/will be used for public data releases from Mayo MCSA/ADRC, ADNI4, SCAN, A4/LEARN, ALLFTD, and others
- Freely available for noncommercial research use: https://www.nitrc.org/projects/mri_reface/
  - New version 0.3.2 with Docker support
- Future work on partial de-facing for cancer applications.

MAYO CLINIC

Thanks

Thank you to all our research participants!

Collaborators:
Walter K. Kremers, Terry M. Therneau, Richard R. Sharp, Marios Savvides, Stephen D. Weigand, Carl M. Prakaashana, Jeffrey L. Gunter, Val J. Lowe, Prashanthi Vemuri, Arvin Arani, Anthony J. Spychalla, Kejal Kantarci, David S. Knopman, Ronald C. Petersen, Clifford R. Jack Jr.

Funding:
NIH grants **R01 AG068206**, U01 AG006786, P50 AG016574, R01 AG034676, R37 AG011378, R01 AG041851, R01 NS097495, R01 AG056366, U01 NS100620; The GHR Foundation; The Elsie and Marvin Dekelboum Family Foundation; The Alexander Family Alzheimer's Disease Research Professorship of the Mayo Clinic; The Liston Award; The Schuler Foundation; and The Mayo Foundation for Medical Education and Research. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of a Quadro P6000 GPU used in generating 3D facial reconstructions for this research. We also thank AVID Radiopharmaceuticals, Inc., for their support in supplying Flortaucipir precursor, chemistry production advice, and FDA regulatory cross-filing permission and documentation needed for this work.

MAYO CLINIC