

# Understanding Sources of Demographic Predictability in Brain MRI via Disentangling Anatomy and Contrast

Mehmet Yigit Avci<sup>\*,†1</sup>, Akshit Acharya<sup>\*1</sup>, Andrew P. King<sup>†1</sup>, and Jorge Cardoso<sup>†1,2</sup>, for the Alzheimer’s Disease Neuroimaging Initiative

King’s College London, UK

{yigit.avci,akshit.acharya,m.jorge.cardoso,andrew.king}@kcl.ac.uk

**Abstract.** Demographic attributes can be predicted from medical images, raising concerns about bias in clinical AI systems. In X-ray imaging, acquisition characteristics have been shown to contribute substantially to this predictability. Whether the same holds in brain MRI remains unclear, as anatomical variation and acquisition-dependent contrast are deeply entangled in the image formation process, obscuring the origins of demographic signal. To address this, we propose a controlled framework based on disentangled representation learning, decomposing brain MRI into anatomy-focused representations that suppress acquisition influence and contrast embeddings that capture acquisition-dependent characteristics. Training predictive models for age, sex, and race on full images, anatomical representations, and contrast embeddings allows us to quantify the relative contributions of structure and acquisition to the demographic signal. Across three datasets and multiple MRI sequences, demographic predictability is found to be driven primarily by anatomical variation, with anatomy-focused representations largely preserving the performance of models trained on raw images. Contrast embeddings retain a weaker signal that is dataset-specific and does not generalize across sites. These findings suggest that effective mitigation must explicitly account for the primarily anatomical and secondarily acquisition-dependent origins of demographic signal, ensuring that any bias reduction generalizes robustly across domains.

**Keywords:** Fairness · Bias · Disentanglement · Representation Learning

## 1 Introduction

Medical imaging is essential for clinical practice and biomedical research, supporting diagnosis, treatment planning, and disease monitoring. While machine learning has enabled powerful image-based models for tasks such as segmentation, classification, and prognosis, growing evidence suggests that these models

---

\* These authors contributed equally to this work.

† These authors jointly supervised this work.

‡ Corresponding author: yigit.avci@kcl.ac.uk.

can exploit unintended correlations in imaging data, leading to performance disparities across demographic groups and allowing attributes such as race to remain identifiable even when unrelated to the clinical task [23,15,26,20,8]. Attempts to mitigate these effects through generic fairness interventions have proven insufficient, often degrading performance or failing under real-world distribution shifts [22]. These observations suggest that the demographic signal is not a superficial modelling artifact but reflects structured variation embedded within the data itself.

Recent work has begun to investigate where this signal originates. For example, in cardiac MRI segmentation, demographic information has been shown to persist within learned representations despite controlled task settings [16]. In chest X-ray classification, acquisition-related factors substantially contribute to the demographic signal and downstream performance disparities [17]. Collectively, these findings indicate that bias may also emerge from modality and acquisition-specific characteristics that models implicitly exploit.

This perspective is particularly relevant for brain MRI. Although demographic bias has been documented in MRI-based models [1,19,7], its origins remain unclear. Brain MRI is intrinsically heterogeneous as image appearance is governed by interactions between sequence parameters (e.g., repetition time (TR), echo time (TE)), scanner hardware, and site-specific protocols. These acquisition-dependent factors systematically shape intensity distributions and tissue contrast, introducing variability sources that correlate with demographic attributes. As anatomical structure and acquisition-dependent contrast are entangled in raw images, conventional models cannot determine whether the demographic signal arises from biological variation, acquisition settings, or their interaction.

To understand the underlying drivers of demographic bias in brain MRI, we employ disentangled representation learning frameworks, specifically MR-CLIP [4] and DIST-CLIP [3]. These models decompose brain MRI into distinct anatomical images and contrast embeddings, enabling explicit separation of structural information from acquisition-dependent signal. We train predictive models for age, sex, and race using (i) full brain MRI images, (ii) anatomy-focused images that suppress acquisition effects, and (iii) contrast embeddings that capture acquisition characteristics while minimizing anatomical content. By comparing predictive performance across these representations, we quantify the relative contributions of anatomical structure and acquisition-induced contrast to the demographic signal. Our contributions are: (1) a controlled framework for disentangling anatomical and acquisition-dependent sources of the demographic signal in brain MRI; (2) evidence, across three datasets and multiple MRI sequences, that anatomical structure is the dominant carrier of demographic information; and (3) showing that contrast embeddings retain only a weak, dataset-specific signal that does not generalize across sites, establishing that effective bias mitigation strategies must primarily address anatomical variation. Depending on site-specific demographics and protocols, acquisition characteristics may also require consideration.

## 2 Methods


**Experimental Design.** Let  $\mathcal{D} = \{(x_i, y_i, a_i)\}_{i=1}^N$  denote a dataset of brain MRI scans  $x_i \in \mathcal{X}$ , demographic attributes  $y_i \in \mathcal{Y}$  (age, sex, race), and acquisition metadata  $a_i \in \mathcal{A}$  (e.g., TR, TE, scanner type, site). Our objective is to determine whether the demographic signal arises primarily from anatomical structure, acquisition-dependent contrast, or their interaction.

In raw MRI, anatomical structure and contrast are inherently entangled within the image space  $\mathcal{X}$ . We therefore seek a representation-level decomposition of each image  $x$  into two components,

$$x \longrightarrow (z_{\text{anat}}, z_{\text{contrast}}), \quad (1)$$

where  $z_{\text{anat}} \in \mathcal{Z}_{\text{anat}}$  captures structural information while minimizing acquisition effects, and  $z_{\text{contrast}} \in \mathcal{Z}_{\text{contrast}}$  captures acquisition-dependent signal while reducing anatomical dominance. To operationalize this separation, we leverage pre-trained disentangled representation learning frameworks, MR-CLIP [4] and DIST-CLIP [3]. MR-CLIP is a metadata-guided contrastive learning framework that aligns MRI images with their associated DICOM acquisition parameters without requiring manual sequence labels. An image encoder  $E_I$  and a metadata encoder  $E_M$  are trained jointly so that the image embedding  $z_{\text{contrast}}^x = E_I(x)$  and the metadata embedding  $z_{\text{metadata}}^a = E_M(a)$  are similar for matched pairs and dissimilar for mismatches. By optimizing a contrastive objective over image-metadata pairs, MR-CLIP learns a 512-dimensional embedding that reflects acquisition differences [2,4]. DIST-CLIP builds upon MR-CLIP by introducing an explicit anatomical pathway to separate structural content from acquisition-dependent contrast. It employs an anatomy mapper that projects an input image  $x$  to an anatomy-focused representation  $z_{\text{anat}}$ . This representation is designed to preserve geometric and morphological structure while suppressing intensity and contrast-related cues. To enforce acquisition invariance, DIST-CLIP incorporates contrastive loss applied at the patch level across anatomy images. This loss encourages corresponding anatomical patches from images acquired with different protocols to remain similar in the  $z_{\text{anat}}$  space, regardless of acquisition differences. This explicitly promotes structural consistency while discouraging the encoding of contrast-specific information. As a result, the model achieves a disentangled decomposition in which  $z_{\text{anat}}$  captures acquisition-invariant information, while  $z_{\text{contrast}}$  encodes acquisition-dependent characteristics [3].

Fig. 1 provides qualitative and quantitative evidence of the representation separation. Panel A shows raw T1-weighted images from OASIS, HCP, and ADNI alongside their anatomy-focused counterparts ( $z_{\text{anat}}$ ): while raw images vary visibly in intensity and tissue contrast across sites, the anatomical representations retain structural detail and suppress these acquisition-driven differences. Panel B visualizes contrast embeddings ( $z_{\text{contrast}}$ ) via t-SNE, where clear clustering by dataset and MRI sequence confirms that acquisition-dependent characteristics are well captured. The framework does not assume perfect disentanglement, an inherently unattainable objective, but instead we explicitly



**Fig. 1.** Qualitative and embedding-level visualization of the disentanglement. **A:** Raw images ( $x$ , top) and anatomy-focused representations ( $z_{\text{anat}}$ , bottom). **B:** t-SNE of contrast embeddings ( $z_{\text{contrast}}$ ) colored by dataset and sequence. **C:** Linear probing ( $N=500$ ). Anatomical probes predict intracranial volume (ICV) and ventricle volume; the acquisition probe predicts scanner/protocol cluster (15 classes; chance = 6.7%).  $z_{\text{metadata}}$  denotes the metadata-only embedding from  $E_M$ .

quantify the degree of disentanglement through a probing analysis (Panel C). Using SynthSeg [5], we extract ground-truth intracranial volume (ICV) and ventricle volume for  $N=500$  OASIS-3 [14] subjects, then train linear regression and classification probes on dimensionally matched representations under Group K-Fold cross-validation. The results show that  $z_{\text{anat}}$  strongly predicts anatomical features (e.g., ICV  $R^2 = 0.953$ , MAE = 31.1 cm<sup>3</sup>) but poorly predicts acquisition protocol (BalAcc 42.5%, where chance is 6.7% over 15 scanner/protocol clusters). Conversely,  $z_{\text{contrast}}$  is informative about acquisition identity (BalAcc 70.2%) while encoding substantially less anatomical detail (ICV  $R^2 = 0.805$ , MAE = 88.1 cm<sup>3</sup>). Crucially, the metadata-only embedding  $z_{\text{metadata}}$ , derived from DICOM parameters with no access to the image, already achieves (ICV  $R^2 = 0.739$ , MAE = 94.2 cm<sup>3</sup>), indicating that most of the residual anatomical signal in  $z_{\text{contrast}}$  does not reflect true image-level leakage, but rather population-level correlations between acquisition settings and demographic composition across sites.

To quantify the contribution of each representation to the demographic signal, we define three predictive mappings:

$$\begin{aligned}
 f_{\text{full}} &: \mathcal{X} \rightarrow \mathcal{Y}, \\
 f_{\text{anat}} &: \mathcal{Z}_{\text{anat}} \rightarrow \mathcal{Y}, \\
 f_{\text{contrast}} &: \mathcal{Z}_{\text{contrast}} \rightarrow \mathcal{Y}.
 \end{aligned}$$

Here,  $f_{\text{full}}$  is trained directly on raw MRI  $x$ , while  $f_{\text{anat}}$  and  $f_{\text{contrast}}$  operate on disentangled representations  $z_{\text{anat}}$  and  $z_{\text{contrast}}$ , respectively. For each

demographic attribute  $y \in \{\text{age, sex, race}\}$ , we train models under identical optimization protocols and compare predictive performance across these mappings. Performance differences between  $f_{\text{anat}}$  and  $f_{\text{contrast}}$  provide an empirical estimate of the relative contribution of anatomical structure and acquisition-dependent contrast to the demographic signal. The performance of  $f_{\text{full}}$  serves as a reference for the total demographic information accessible from the raw image  $x$ .

**Models and Evaluation.** For image-based inputs ( $f_{\text{full}}$  and  $f_{\text{anat}}$ ), we employ a 3D ResNet-50 architecture [10,24] operating on volumetric inputs of size  $128 \times 128 \times 128$ . All models are initialized from ImageNet-pretrained weights (inflated to 3D). A fixed learning rate of  $1 \times 10^{-4}$  is used across all experiments with dropout regularization of 0.2 to mitigate overfitting. The batch size is 2 due to memory constraints of 3D training. For embedding-based inputs ( $f_{\text{contrast}}$ ), we use a multi-layer perceptron operating on  $z_{\text{contrast}}$  with architecture  $512 (\text{input}) \rightarrow 512 \rightarrow 256 \rightarrow N_{\text{classes}}$ , with ReLU activations and dropout (0.2) applied after each hidden layer.

Age prediction is formulated as a regression task optimized using mean squared error (MSE) (HCP age uses range midpoint). Sex and race prediction are formulated as classification tasks optimized with cross-entropy loss. All models are trained under matched data splits, learning rates, and regularization schemes to ensure comparability across representation types. We report mean absolute error (MAE) for age prediction and balanced accuracy (BalAcc) for classification tasks, ensuring equal weighting of classes to account for potential imbalance. Balanced accuracy is defined as the average recall across classes (i.e., mean per-class sensitivity), thereby preventing inflated performance estimates in the presence of class imbalance.

**Datasets and Preprocessing.** We evaluate our approach on three publicly available brain MRI datasets: Open Access Series of Imaging Studies (OASIS) [14], Alzheimer’s Disease Neuroimaging Initiative (ADNI) [12], and Human Connectome Project (HCP) [25]. These cohorts span different age ranges, scanner vendors, and acquisition protocols, providing diverse combinations of  $(x_i, y_i, a_i)$ . This heterogeneity is essential for disentangling anatomical structure from contrast variations induced by imaging parameters. The demographic distribution of each dataset, including age, sex, and race, is summarized in Table 1. To ensure consistency across datasets and compatibility with the disentanglement framework, all volumes are rigidly registered to the MNI152 atlas at  $1.0 \text{ mm}^3$  isotropic resolution. Skull stripping is performed using SynthStrip [11].

### 3 Experiments and Results

To examine how anatomical and acquisition-dependent information contribute to demographic predictability, we compare the signal available from each representation within individual datasets and under joint training, test whether this predictability transfers across datasets as cohort and acquisition characteristics change, and assess whether the same anatomy–contrast pattern holds across MRI sequences.

**Table 1.** Subject-level demographic summary of the included datasets. Age is reported as mean  $\pm$  standard deviation. Sex and race are reported as counts (percentage within dataset).

| Dataset | Age            | Sex         |             | Race         |             |
|---------|----------------|-------------|-------------|--------------|-------------|
|         | Mean $\pm$ SD  | Female      | Male        | White        | Black       |
| OASIS-3 | 72.4 $\pm$ 9.1 | 608 (55.6%) | 485 (44.4%) | 926 (84.7%)  | 167 (15.3%) |
| ADNI    | 74.7 $\pm$ 7.3 | 725 (44.6%) | 901 (55.4%) | 1542 (94.8%) | 84 (5.2%)   |
| HCP     | 29.1 $\pm$ 3.5 | 543 (54.4%) | 455 (45.6%) | 830 (83.2%)  | 168 (16.8%) |

The within-dataset and joint-training results for T1-weighted images are summarized in Table 2. Models trained on raw images achieve the highest overall performance across tasks (e.g., sex BalAcc 0.92–0.96; age MAE 2.67–5.15 years). Anatomy-focused representations largely preserve this performance, for instance, sex prediction reaches 0.93 in OASIS, indicating that substantial demographic signal is embedded in structural variation. When training jointly across all datasets, this pattern remains consistent (Sex 0.96 Raw vs. 0.96 Anat vs. 0.78 Contrast), suggesting that the dominance of anatomical signal is not driven by single-dataset effects. Importantly, contrast embeddings, despite explicitly minimizing anatomical content, retain non-trivial predictive power. Sex prediction from  $z_{\text{contrast}}$  achieves BalAcc 0.71 in OASIS and 0.81 in HCP (0.78 in joint training), while race prediction remains above chance in OASIS (0.71), and joint training (0.68). Even for age, contrast embeddings yield informative MAE values (e.g., 6.28 years in OASIS, and 5.21 jointly). Although consistently lower than those of the raw-image and anatomy-based models, these results demonstrate that contrast alone retains a demographic signal when evaluated within a single domain.

**Table 2.** Within-dataset and joint-training performance across representations. For each dataset (OASIS, ADNI, HCP) and joint training across all datasets (ALL), we report balanced accuracy (BalAcc,  $\uparrow$ ; equal class weighting) for **Sex** and **Race**, and mean absolute error (MAE,  $\downarrow$ ) for **Age**.

| Dataset | Sex (BalAcc $\uparrow$ ) |      |          | Race (BalAcc $\uparrow$ ) |      |          | Age (MAE $\downarrow$ ) |      |          |
|---------|--------------------------|------|----------|---------------------------|------|----------|-------------------------|------|----------|
|         | Raw                      | Anat | Contrast | Raw                       | Anat | Contrast | Raw                     | Anat | Contrast |
| OASIS   | 0.92                     | 0.93 | 0.71     | 0.86                      | 0.81 | 0.71     | 5.15                    | 5.45 | 6.28     |
| ADNI    | 0.94                     | 0.93 | 0.74     | 0.97                      | 0.82 | 0.50     | 4.46                    | 4.43 | 5.61     |
| HCP     | 0.95                     | 0.96 | 0.81     | 0.96                      | 0.90 | 0.70     | 2.67                    | 2.99 | 2.77     |
| ALL     | 0.96                     | 0.96 | 0.78     | 0.91                      | 0.93 | 0.68     | 3.91                    | 4.35 | 5.21     |

**Table 3.** Cross-dataset generalization. **Sex** and **Race** are reported as balanced accuracy (BalAcc,  $\uparrow$ ; equal class weighting), and **Age** is reported as mean absolute error (MAE,  $\downarrow$ ).

| Train | Test  | Sex (BalAcc $\uparrow$ ) |      |          | Race (BalAcc $\uparrow$ ) |      |          | Age (MAE $\downarrow$ ) |       |          |
|-------|-------|--------------------------|------|----------|---------------------------|------|----------|-------------------------|-------|----------|
|       |       | Raw                      | Anat | Contrast | Raw                       | Anat | Contrast | Raw                     | Anat  | Contrast |
| OASIS | HCP   | 0.65                     | 0.82 | 0.50     | 0.85                      | 0.87 | 0.50     | 26.37                   | 28.03 | 29.89    |
|       | ADNI  | 0.88                     | 0.89 | 0.59     | 0.93                      | 0.99 | 0.53     | 5.15                    | 5.50  | 8.71     |
| HCP   | OASIS | 0.72                     | 0.71 | 0.53     | 0.85                      | 0.82 | 0.70     | 39.55                   | 42.42 | 46.91    |
|       | ADNI  | 0.60                     | 0.61 | 0.50     | 0.88                      | 0.84 | 0.60     | 42.39                   | 44.24 | 50.61    |
| ADNI  | OASIS | 0.76                     | 0.84 | 0.70     | 0.81                      | 0.84 | 0.50     | 5.17                    | 5.49  | 6.47     |
|       | HCP   | 0.62                     | 0.81 | 0.51     | 0.86                      | 0.90 | 0.50     | 29.58                   | 28.11 | 30.77    |

To determine whether demographic signals reflect biological variation or site-related acquisition-specific effects, we evaluate cross-dataset generalization by training on one dataset and testing on another. This introduces shifts in acquisition protocols, scanner vendors, and population characteristics, allowing us to assess whether demographic signals persist under distribution shift (Table 3). Across all tasks, performance drops markedly under dataset shift, particularly for age prediction, where MAE increases dramatically (e.g., HCP $\rightarrow$ ADNI: Raw 42.39 years). However, age-transfer results involving HCP should be interpreted in the context of the markedly different cohort age ranges. HCP participants span 23.5–36 years, whereas OASIS and ADNI participants span 42–97 and 55–95 years, respectively. These experiments therefore capture both acquisition-domain shift and a substantial shift in the age distribution. For classification, a consistent pattern emerges: anatomy-focused representations ( $z_{\text{anat}}$ ) tend to generalize more robustly than raw images. For example, ADNI $\rightarrow$ OASIS sex prediction improves from BalAcc 0.76 (Raw) to 0.84 (Anat). Whereas, contrast embeddings ( $z_{\text{contrast}}$ ) frequently collapse toward chance (0.50), indicating poor transferability of acquisition-dependent signal. These findings suggest that while contrast embeddings encode measurable demographic information within a dataset, much of this signal is dataset-specific and does not generalize across sites. Structural representations, in contrast, retain more stable demographic information under distribution shift. Together with the within-dataset results, this indicates that demographic predictability in brain MRI arises mainly from anatomical sources, and acquisition-driven correlations are less robust and more site-dependent.

We further investigate the demographic signal across imaging sequences within OASIS (Table 4) to assess whether the relative roles of anatomy and contrast vary with sequence type. Across T2w and FLAIR images, anatomy representations remain close to raw performance (e.g., sex BalAcc 0.91 vs. 0.90 for T2w), indicating that structural variation continues to account for a substantial portion of the demographic signal beyond T1-weighted imaging. Contrast embeddings retain measurable predictive power across sequences. For example, sex prediction from  $z_{\text{contrast}}$  reaches BalAcc 0.72 in T2w and 0.75 in FLAIR whereas race prediction remains near chance. Joint training on all sequences maintains strong

sex prediction performance, suggesting that combining contrasts does not remove demographic signal. Anatomy representations continue to perform strongly (sex BalAcc 0.93, age MAE 5.16), while contrast embeddings remain weaker yet informative. Overall, these results show that demographic predictability extends beyond T1-weighted imaging where anatomy-focused representations largely contain the demographic signal across sequences and contrast embeddings contain a weaker, more sequence-dependent signal.

## 4 Discussion

In this study, we go beyond showing that demographic attributes can be predicted from brain MRI to address a more fundamental question: where does this signal originate? By decomposing MRI into anatomy-focused and contrast-dependent representations, our framework examines the relative contributions of anatomical structure and acquisition-dependent contrast to demographic predictability.

Our results consistently show that the majority of the demographic signal is preserved in anatomy-focused representations, establishing structural variation as the dominant contributor. This aligns with biological expectations for age and is plausible for sex, given known morphological differences [18,9,21], and there is evidence that relative brain volume can differ across racial and ethnic groups [6,13]. At the same time, contrast embeddings retain stable, above-chance predictive performance across datasets and sequences which might be due to population-level acquisition correlations or minor disentanglement leakage.

These findings carry important implications. Because a substantial portion of the demographic signal resides in anatomical structure, approaches that focus exclusively on acquisition harmonization or intensity normalization are unlikely to fully mitigate demographic predictability in downstream applications such as Alzheimer’s disease detection. Conversely, strategies that indiscriminately remove anatomical variation risk suppressing clinically meaningful information. Effective bias mitigation should therefore focus on identifying and controlling the specific anatomical features that contribute to demographic signal while preserv-

**Table 4.** OASIS multi-sequence analysis. Performance across sequences for models trained on raw images ( $f_{\text{full}}$ ), anatomy-focused representations ( $f_{\text{anat}}$ ), and contrast-only embeddings ( $f_{\text{contrast}}$ ). **Sex** and **Race** are reported as balanced accuracy (BalAcc,  $\uparrow$ ; equal class weighting). **Age** is reported as mean absolute error (MAE,  $\downarrow$ ).

| Sequence      | Sex (BalAcc $\uparrow$ ) |      |          | Race (BalAcc $\uparrow$ ) |      |          | Age (MAE $\downarrow$ ) |      |          |
|---------------|--------------------------|------|----------|---------------------------|------|----------|-------------------------|------|----------|
|               | Raw                      | Anat | Contrast | Raw                       | Anat | Contrast | Raw                     | Anat | Contrast |
| T1w           | 0.92                     | 0.93 | 0.71     | 0.86                      | 0.81 | 0.71     | 5.15                    | 5.45 | 6.28     |
| T2w           | 0.91                     | 0.90 | 0.72     | 0.74                      | 0.88 | 0.51     | 5.44                    | 5.91 | 6.60     |
| FLAIR         | 0.71                     | 0.83 | 0.75     | 0.65                      | 0.85 | 0.50     | 4.54                    | 4.80 | 6.67     |
| All Sequences | 0.91                     | 0.93 | 0.76     | 0.84                      | 0.81 | 0.70     | 5.37                    | 5.16 | 5.90     |

ing those relevant to the clinical task. Acquisition-related variations should be addressed when they introduce additional demographic information, particularly in site-specific settings.

Identifying demographic signal does not in itself imply harmful bias in clinical predictions. However, understanding where this signal resides is a prerequisite for evaluating its downstream impact. Our framework provides a basis for investigating how these components propagate into clinical tasks and under what conditions they may translate into measurable bias, supporting the development of causally informed and fair neuroimaging models.

#### 4.1 Limitations

The image-based predictors are based on 3D ResNet-50, whereas the contrast embeddings are evaluated using a smaller MLP because the inputs differ in dimensionality and structure. Consequently, performance differences may partly reflect predictor capacity and should not be interpreted as a pure measure of the information contained in each representation. Analysis on the extent to which this affects the interpretations is left for future work.

Another limitation of this study is the imbalanced racial composition of the included datasets, particularly ADNI, where only 84 participants (5.2%) are Black. Although balanced accuracy mitigates the effect of class imbalance on the reported metrics, the small minority-class size limits the statistical reliability of race prediction results. Future work should evaluate these effects in more demographically representative cohorts.

**Acknowledgments.** This research was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) [grant reference number EP/Y035216/1] Centre for Doctoral Training in Data-Driven Health (DRIVE-Health) at King’s College London, with additional support from deepc GmbH. Further support was provided by the Scientific and Technological Research Council of Türkiye (TÜBİTAK) under the 2213-A Overseas Graduate Scholarship. Data used in preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (<https://adni.loni.usc.edu>). As such, the investigators within ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in the analysis or writing of this report. A complete listing of ADNI investigators can be found at [https://adni.loni.usc.edu/wp-content/uploads/how\\_to\\_apply/ADNI\\_Acknowledgement\\_List.pdf](https://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf). Data were also provided in part by OASIS and the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657), funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research, and by the McDonnell Center for Systems Neuroscience at Washington University.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Achara, A., Anton, E.P., Hammers, A., King, A.P.: Invisible attributes, visible biases: Exploring demographic shortcuts in mri-based alzheimer’s disease classification (2025), <https://arxiv.org/abs/2509.09558>
2. Avci, M.Y., Borges, P., Wright, P., Yigitsoy, M., Ourselin, S., Cardoso, J.: MR-CLIP: Efficient metadata-guided learning of MRI contrast representations. In: Cui, Z., Rekik, A., Suk, H.I., Ouyang, X., Sun, K., Wang, S. (eds.) *Machine Learning in Medical Imaging. MLMI 2025, Lecture Notes in Computer Science*, vol. 16241. Springer, Cham (2026)
3. Avci, M.Y., Borges, P., Fernandez, V., Wright, P., Yigitsoy, M., Ourselin, S., Cardoso, J.: DIST-CLIP: Arbitrary metadata and image guided MRI harmonization via disentangled anatomy-contrast representations (2025), <https://arxiv.org/abs/2512.07674>
4. Avci, M.Y., Borges, P., Fernandez, V., Wright, P., Yigitsoy, M., Ourselin, S., Cardoso, J.: Metadata-aligned 3D MRI representations for contrast understanding and quality control. *arXiv:2511.00681* (2025)
5. Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Leemput, K.V., Fischl, B., Dalca, A.V., Iglesias, J.E.: SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical Image Analysis* **86**, 102789 (2023)
6. Brickman, A.M., Schupf, N., Manly, J.J., Luchsinger, J.A., Andrews, H., Tang, M.X., Reitz, C., Small, S.A., Mayeux, R., DeCarli, C., Brown, T.R.: Brain morphology in older african americans, caribbean hispanics, and whites from northern manhattan. *Archives of Neurology* **65**(8), 1053–1061 (08 2008)
7. Danaee, G., Niethammer, M., Rushmore, J., Bouix, S.: Investigating demographic bias in brain mri segmentation: A comparative study of deep-learning and non-deep-learning methods. *Machine Learning for Biomedical Imaging* **3**, 792–808 (2025)
8. Gichoya, J.W., Banerjee, I., Bhimireddy, A.R., Burns, J.L., Celi, L.A., Chen, L.C., Correa, R., Dullerud, N., Ghassemi, M., Huang, S.C., Kuo, P.C., Lungren, M.P., Palmer, L.J., Price, B.J., Purkayastha, S., Pyrros, A.T., Oakden-Rayner, L., Okechukwu, C., Seyyed-Kalantari, L., Trivedi, H., Wang, R., Zaiman, Z., Zhang, H.: AI recognition of patient race in medical imaging: a modelling study. *The Lancet Digital Health* **4**(6), e406–e414 (2022)
9. Good, C.D., Johnsrude, I., Ashburner, J., Henson, R.N., Friston, K.J., Frackowiak, R.S.: Cerebral asymmetry and the effects of sex and handedness on brain structure: a voxel-based morphometric analysis of 465 normal adult human brains. *Neuroimage* **14**(3), 685–700 (2001)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition (2015), <https://arxiv.org/abs/1512.03385>
11. Hoopes, A., Mora, J.S., Dalca, A.V., Fischl, B., Hoffmann, M.: SynthStrip: skull-stripping for any brain image. *NeuroImage* **260**, 119474 (2022)
12. Jack, C.R.J., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., Whitwell, J.L., Ward, C., Dale, A.M., Felmlee, J.P., Gunter, J.L., Hill, D.L., Killiany, R., Schuff, N., Fox-Bosetti, S., Lin, C.P., Studholme, C., DeCarli, C.S., Krueger, G., Ward, H.A., Metzger, G.J., Scott, K.T., Mallozzi, R., Blezek, D., Levy, J., Debbins, J.P., Fleisher, A.S., Albert, M., Green, R., Bartzokis, G., Glover, G., Mugler, J., Weiner, M.W.: The Alzheimer’s Disease Neuroimaging Initiative (ADNI): MRI methods. *Journal of Magnetic Resonance Imaging* **27**(4), 685–691 (2008)

13. Kang, D.W., Wang, S.M., Na, H.R., Park, S.Y., Kim, N.Y., Lee, C.U., Kim, D., Son, S.J., Lim, H.K.: Differences in cortical structure between cognitively normal east asian and caucasian older adults: a surface-based morphometry study. *Scientific reports* **10**(1), 20905 (2020)
14. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A., Raichle, M.E., Cruchaga, C., Marcus, D.: Oasis-3: Longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *medRxiv* (2019)
15. Larrazabal, A.J., Nieto, N., Peterson, V., Milone, D.H., Ferrante, E.: Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. *Proceedings of the NAS* **117**(23), 12592–12594 (2020)
16. Lee, T., Puyol-Anton, E., Ruijsink, B., Masci, P.G., Keehn, L., Chowienzyk, P., Haseler, E., Shi, M., King, A.: Understanding-informed bias mitigation for fair CMR segmentation. *Machine Learning for Biomedical Imaging* **3**, 808–824 (2025)
17. Lotter, W.: Acquisition parameters influence AI recognition of race in chest X-rays and mitigating these factors reduces underdiagnosis bias. *Nature Communications* **15**, 7465 (2024)
18. Passe, T.J., Rajagopalan, P., Tupler, L.A., Byrum, C.E., Macfall, J.R., Krishnan, K.R.: Age and sex effects on brain morphology. *Progress in Neuro-Psychopharmacology and Biological Psychiatry* **21**(8), 1231–1237 (1997)
19. Petersen, E., Feragen, A., da Costa Zemsch, M.L., Henriksen, A., Wiese Christensen, O.E., Ganz, M.: Feature robustness and sex differences in medical imaging: A case study in MRI-based Alzheimer’s Disease detection. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*. pp. 88–98. Springer Nature Switzerland, Cham (2022)
20. Puyol-Antón, E., Ruijsink, B., Piechnik, S.K., Neubauer, S., Petersen, S.E., Razavi, R., King, A.P.: Fairness in cardiac mr image analysis: An investigation of bias due to data imbalance in deep learning based segmentation. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 413–423. Springer International Publishing, Cham (2021)
21. Ruigrok, A.N., Salimi-Khorshidi, G., Lai, M.C., Baron-Cohen, S., Lombardo, M.V., Tait, R.J., Suckling, J.: A meta-analysis of sex differences in human brain structure. *Neuroscience & Biobehavioral Reviews* **39**, 34–50 (2014)
22. Schrouff, J., Harris, N., Koyejo, O., Alabdulmohsin, I., Schnider, E., Opsahl-Ong, K., Brown, A., Roy, S., Mincu, D., Chen, C., Dieng, A., Liu, Y., Natarajan, V., Karthikesalingam, A., Heller, K., Chiappa, S., D’Amour, A.: Diagnosing failures of fairness transfer across distribution shift in real-world medical settings (2023), <https://arxiv.org/abs/2202.01034>
23. Seyyed-Kalantari, L., Zhang, H., McDermott, M., Chen, I., Marzyeh, G.: Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. *Nature Medicine* **27**, 2176–2182 (2021)
24. Solovyev, R., Kalinin, A.A., Gabruseva, T.: 3d convolutional neural networks for stalled brain capillary detection. *Computers in Biology and Medicine* **141**, 105089 (2022)
25. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K.: The WU-Minn Human Connectome Project: An overview. *NeuroImage* **80**, 62–79 (2013)
26. Waite, S., Scott, J., Colombo, D.: Narrowing the gap: Imaging disparities in radiology. *Radiology* **299**(1), 27–35 (2021), PMID: 33560191