# Using MR physics for domain generalisation and super-resolution

Pedro Borges[1], Virginia Fernandez[1], Petru Daniel Tudosiu[1], Parashkev Nachev[2], Sebastien Ourselin[1], and M. Jorge Cardoso[1]

[1]School of Biomedical Engineering and Imaging Sciences, KCL, UK
[2]Queen Square Institute of Neurology, UCL, UK

**Abstract.** MRI is a very flexible imaging modality, but with flexibility comes heterogeneity. MRI sequence choice, acquisition parameters, and image resolution form an extrinsic source of variability, reducing our ability to extract the underlying relevant biological signal and causing difficulties in downstream analyses. We propose a new method that can create resolution and acquisition-parameter invariant representations by removing external sources of variability. We use realistic physics models of image resolution and combine them with a differentiable model of MRI sequences to create an invariant high-resolution multi-parametric (MPM) MRI estimate from an arbitrary number of inputs, all trained via self-supervision. The proposed method allows clinical imaging sessions with sequences acquired at arbitrary resolutions to be transformed into a single-domain generalisable representation. We demonstrate the model's validity by showing improved MPM reconstruction and imputation quality compared to previous methods and a significantly improved ability to super-resolve. We also demonstrate domain generalisation capabilities via a downstream classification model that is more robust to the choice of input sequences in an out-of-distribution dataset.

## 1 Introduction

MRI is a versatile imaging modality that excels at providing soft tissue contrast. The sequence choice highlights different tissue properties depending on the task. Owing to each MR sequence's bespoke nature, images acquired in routine clinical care present widely varying resolutions and contrasts, beholden to the choice of sequence. The joint processing of images arising from such clinical sources is fraught with challenges, as algorithms typically favour homogeneity [15].

Variability in resolution can be mitigated via super-resolution (SR); SR methods strive to increase image resolution by leveraging complementary information within and across multiple subject images and have been shown to approximate acquiring images in high-resolution [13].

Quantitative imaging can address the contrast-variability aspect, as values become individually meaningful, and information can be derived directly rather than through a qualitative surrogate. However, quantitative maps (multiparametric maps, or MPMs) are rare in a clinical setting due to the acquisition's

complex and time-consuming nature (requires multi-echo data). Rather than acquiring optimized data for MPMs, many methods propose to generate these MPMs from fewer, often arbitrary, data points [10, 16, 8, 1]. In this work, we build upon Borges et al.'s method (we denote it MPMGen), due to its least restrictive data needs, purporting to be able to generate MPMs self-supervised and in the absence of some modalities. Note that while all these handle sequence variability, no method currently does so while modelling image resolution, neither implicitly nor explicitly. This is particularly important because MRI images from different sequences are normally acquired with different resolutions.

We propose SR-MPMGen, significantly building upon MPMGen, by enabling joint modelling acquisition physics and resolution. SR-MPMGen is, therefore, capable of addressing the heterogeneity introduced by varying contrast (MPM-conversion) and resolution (super-resolution), allowing for the creation of a generalisable and invariant representation of MRI data.

## 2 Methods

### 2.1 Building on a foundation of contrast generalisation

The MPMGen model [1] can convert data acquired from standard clinical pipelines (qualitative data) into an invariant quantitative domain. The work argued that this would allow one to dispense with post hoc domain adaptation techniques, as data translated to such an invariant domain would not require them.

MPMGen accepts any number of volumetric images acquired from MPRAGE, FLAIR, and spin echo (SE) sequences, supporting omissions of up to two of these. Via the use of a multimodal self-attention mechanism (MMSA) [9] and cross-attention (to allow the inclusion of image-specific acquisition parameters) produced a four-dimensional output, meant to contain the quantitative $T_1$, $T_2$, and $PD$ maps (MPMs) for said subject. The output was encouraged to take this form because of a crucial physics-based simulation step, where the output was passed to a series of simplified Bloch equations (one for each sequence) alongside the acquisition parameters of the original images to reproduce the exact input images. An L1 loss between these reconstructions and the original images drove the model's training, with a lower loss indicating higher reconstruction fidelity and, therefore, implicitly, higher MPM fidelity. Notably, the model never sees ground truth MPMs, relying only on commonplace, qualitative images, which are far more readily available.

However, MPMGen omits a crucial component: resolution. A truly invariant representation of MR data should be quantitative and consistent in resolution. Resolution is just as heterogeneous as contrast in clinical MR, as scanning time is precious and isotropy isn't crucial for many diagnostic or exploratory tasks. Therefore, it is common for multi-modal images of any subject to exhibit distinct resolutions. This work uses the original MPMGen architecture but wraps it around a novel super-resolution component and a series of modifications in the training approach to allow the translation of a set of input images into an

isotropic, high-resolution, quantitative space, regardless of the original resolution and acquisition parameters. The proposed model is depicted in Fig. 1.

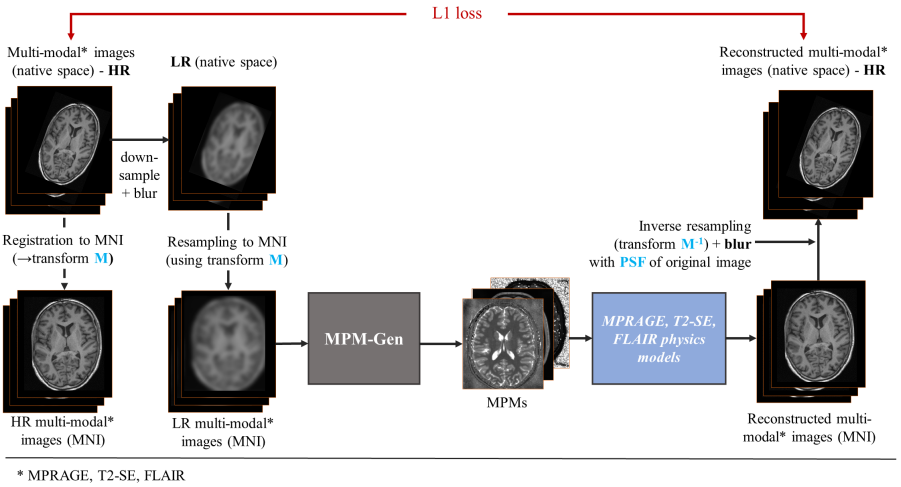


**Fig. 1.** Architecture of SR-MPMGen. Multi-modal images are downsampled and resampled to template space, fed to the MPMGen pipeline to produce MPMs, passed to physics models and resampled to reproduce the original, high-resolution images.

### 2.2 Super-resolution

Following the self-supervised footsteps of the core methodology, we take a similar approach for SR. Downsampling can artificially create low— and high-resolution data. This is not uncommon in super-resolution methods; our approach is inspired by SynthSR [7] but adapted to the multi-sequence MPMGen approach.

We begin by finding a set of affine transformations $M$ registering all subject's input data to 1mm isotropic MNI space [12]. To simulate different resolutions, a downsampling process is designed to produce realistic low-resolution data: slice thickness is mimicked via blurring along the desired axis with a modified Gaussian kernel with a standard deviation proportional to the original and target resolution, given by the following $\sigma = 2\alpha \log(10) \frac{r_t}{r_o} \frac{1}{2\pi}$, where $\alpha = U(0.8, 1.2)$, i.e. drawn from a uniform distribution, $r_t$ is the chosen target axis's new, downsampled, voxel size and $r_o$ is its original resolution. Downsampling is done on the high-resolution images in their native space, then resampling to MNI space using the previously computed transformation $M$. Our target downsampled resolution is selected uniformly between two and seven mm.

The downsampled and MNI resampled images are passed to the core MPM-translation pipeline, producing an MPM-like output. As in the original method, the generated MPM is used in tandem with static equations and the acquisition

parameters for each relevant sequence to reconstruct the original modalities. To allow for direct comparisons, the images are resampled back to their native spaces, but only after correcting for the system's point-spread function (PSF). Applying the PSF is crucial, as resampling to a lower-dimensional grid (such as moving from the isotropic 1x1x1 high-resolution MNI space to a native space of, e.g., 1x1x3 image) would introduce aliasing [3]. In its absence, a high-resolution MPM would be penalised by the loss when presented with an isotropic image (the network produces blurry MPMs), as the resampled image would not match the PSF of the original native image.

### 2.3 Modelling the MR signal via static equations

When applied with MPMGen, the proposed SR method generates high-resolution MPM maps that can be fed into an MRI simulator to regenerate the original non-quantitative data. As in the original work, the MR signal is approximated via a temporally static, computationally efficient model termed static equations, where an entire MR sequence is summarised by a single equation that depends only on the voxelwise intrinsic MR properties and the sequence parameters, fitting well into live deep-learning frameworks. Due to its simple, single-equation formulation, this approach is differentiable, meaning that a simple L1 loss between the original and reconstructed data can drive the SR-MPMGen model training. As an example, the following static equation formulation is used to generate MPRAGE [4], data from the predicted MPMs, for the signal at a single voxel at location $x$:

$$b_{\mathrm{MPRAGE}}(x) = G \cdot PD(x) \left(1 - \frac{2e^{\frac{-TI}{T_1(x)}}}{1 + e^{\frac{-TR}{T_1(x)}}}\right),$$

Here, $TR$ represents repetition time, $TI$ inversion time, and $TE$ echo time. $G$ is a dimensionless parameter that models linear scanner gain. The equivalent equations for FLAIR and T2 spin echo (SE) can be found in [6].

### 2.4 Pre-processing and Data

We employ data from SABRE (932 subjects) [14] and an internal private dataset (IPD) for training our main model. The BraTS dataset [11] is employed for out-of-distribution validation on a downstream task. The inclusion of IPD is motivated by Borges et al.'s model's limited performance on T2-SE-related imputations. We posit that this was due, in part, to the absence of T2-SE images in one of their chosen training datasets, UKB. IPD boasts MPRAGE, FLAIR, and SE images from 837 subjects with great heterogeneity, contrast and resolution-wise, which should be a boon for an algorithm that relies on aptly generalising to a wide range of acquisition types. SABRE likewise contains images of all three sequence types.

## 3 Experiments and results

**Consistency and imputation:** We validate if the translated MPMs suffer due to the introduction of the proposed SR model when compared to vanilla MPMGen. As such, we carry out consistency and imputation analyses in much the same way as the original work. Briefly, consistency analyses aim to verify the fidelity of modality reconstructions where the model has seen said modality. E.g.: MPRAGE + FLAIR → MPM → MPRAGE'. Imputation analyses seek to identify the fidelity of modality reconstructions where the model has not seen said modality. E.g.: MPRAGE + FLAIR → MPM → T2SE'. As a baseline, we train vanilla MPMGen, *MPMGen (New Data)* with the updated datasets. Table 1 and Table 2 show the consistency and imputation results, respectively.

The original paper's MPMGen model boasted higher consistency values ($>$ 0.99 MS-SSIM) while at the same time presenting several imputation failures (MS-SSIM $<$ 0.55), especially where T2-SE reconstructions were involved. *MP-*

**Table 1.** Consistency: Mean MSE and MS-SSIM, for all combinations of MRI modalities, including at least the reconstructed modality on the SABRE dataset. •: Modality present, ∘: Modality absent. MSE is scaled by $10^3$.

| Regenerated modality | Input modalities | | | MPMGen (New Data) | | SR-MPMGen | |
|---|---|---|---|---|---|---|---|
| | FLAIR | MPRAGE | T2-SE | MSE | MS-SSIM | MSE | MS-SSIM |
| FLAIR | • | ∘ | ∘ | $8.57_{15.57}$ | $0.8748_{0.0993}$ | $3.49_{9.58}$ | $0.8953_{0.0535}$ |
| | • | • | ∘ | $7.62_{14.37}$ | $0.878_{0.0921}$ | $3.36_{8.06}$ | $0.8953_{0.0512}$ |
| | • | ∘ | • | $2.48_{0.68}$ | $0.8911_{0.0336}$ | $2.62_{0.85}$ | $0.8857_{0.0348}$ |
| | • | • | • | $2.42_{0.68}$ | $0.8935_{0.0335}$ | $2.59_{0.84}$ | $0.8879_{0.0345}$ |
| MPRAGE | ∘ | • | ∘ | $4.97_{2.32}$ | $0.8624_{0.0512}$ | $4.92_{2.32}$ | $0.8688_{0.0509}$ |
| | • | • | ∘ | $4.18_{1.84}$ | $0.8758_{0.0560}$ | $4.96_{2.30}$ | $0.8660_{0.0500}$ |
| | ∘ | • | • | $4.33_{1.77}$ | $0.8659_{0.0518}$ | $5.10_{2.26}$ | $0.8603_{0.0453}$ |
| | • | • | • | $4.47_{1.83}$ | $0.8611_{0.0484}$ | $5.11_{2.27}$ | $0.8594_{0.0452}$ |
| T2-SE | ∘ | ∘ | • | $2.57_{0.61}$ | $0.8894_{0.0302}$ | $2.14_{0.67}$ | $0.8898_{0.0315}$ |
| | • | ∘ | • | $2.41_{0.54}$ | $0.8923_{0.0290}$ | $2.13_{0.57}$ | $0.8899_{0.0315}$ |
| | ∘ | • | • | $2.46_{0.58}$ | $0.8944_{0.0288}$ | $2.13_{0.57}$ | $0.8915_{0.0325}$ |
| | • | • | • | $2.40_{0.55}$ | $0.8934_{0.0289}$ | $2.14_{0.56}$ | $0.8909_{0.0312}$ |

**Table 2.** Imputation: Mean MSE and MS-SSIM, for all combinations of MRI modalities that do not feature reconstructed modality on the SABRE dataset. •: Modality present, ∘: Modality absent. MSE has been multiplied by $10^3$.

| Regenerated modality | Input modalities | | | MPMGen (New Data) | | SR-MPMGen | |
|---|---|---|---|---|---|---|---|
| | FLAIR | MPRAGE | T2-SE | MSE | MS-SSIM | MSE | MS-SSIM |
| FLAIR | ∘ | • | ∘ | $3.33_{0.65}$ | $0.8684_{0.0339}$ | $3.39_{0.85}$ | $0.86322_{0.03521}$ |
| | ∘ | ∘ | • | $3.35_{0.80}$ | $0.8664_{0.0361}$ | $3.32_{0.81}$ | $0.86793_{0.03633}$ |
| | ∘ | • | • | $2.93_{0.67}$ | $0.8820_{0.0335}$ | $3.00_{0.81}$ | $0.87752_{0.03353}$ |
| MPRAGE | • | ∘ | ∘ | $6.03_{1.89}$ | $0.7935_{0.0556}$ | $6.77_{2.50}$ | $0.79816_{0.05167}$ |
| | ∘ | ∘ | • | $6.15_{2.25}$ | $0.7899_{0.0712}$ | $6.61_{2.63}$ | $0.80459_{0.05703}$ |
| | • | ∘ | • | $5.47_{1.88}$ | $0.8171_{0.0510}$ | $6.36_{2.56}$ | $0.81328_{0.04889}$ |
| T2-SE | • | ∘ | ∘ | $3.31_{0.57}$ | $0.8829_{0.0287}$ | $2.93_{0.59}$ | $0.88762_{0.02616}$ |
| | ∘ | • | ∘ | $4.15_{0.69}$ | $0.8662_{0.0272}$ | $3.67_{0.87}$ | $0.87067_{0.02849}$ |
| | • | • | ∘ | $2.93_{0.57}$ | $0.8931_{0.0283}$ | $2.70_{0.55}$ | $0.89271_{0.02666}$ |

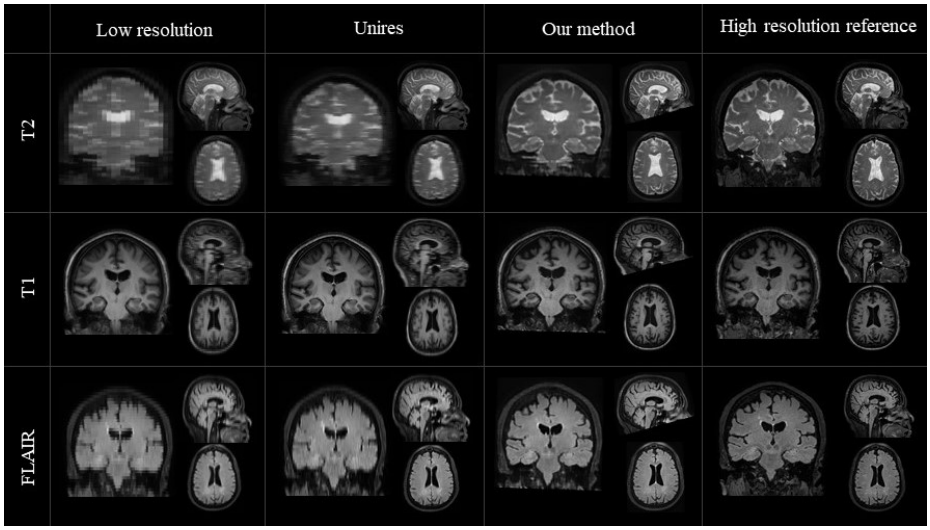


**Fig. 2.** Left-to-right: Three randomly downsampled input modalities, super-resolved images using Unires, super-resolved MPM-derived reconstructions with the proposed method, and the high-resolution reference. Coronal, axial, and sagittal views are shown.

*MGen (New Data)* exhibits lower consistency results across the board alongside more consistent imputation results (No failures). This suggests that the original MPMGen was overfitting to specific modalities given the disparity between results and failure cases, while the similarity between SR-MPMGen's consistency and imputation metrics indicates the opposite, greater robustness and ability to generalise. This behaviour is partially attributed to introducing the highly heterogeneous IPD dataset, which replaced the highly homogeneous Biobank dataset, affording the model a far more diverse set of multi-sequence and parameter images to learn from. Crucially, we observe that SR-MPMGen performs comparably to *MPMGen (New Data)*, indicating that the SR components are not deleterious for MPM translation.

**Super-resolution:** To evaluate the super-resolving abilities of our algorithm, we carry out the downsampling procedure on our data, pass them to our model, and assess the similarity of the physics-reconstructed output to the original, higher-resolution image. The downsampling process is as follows: an axis is selected (at random, for isotropic images, or the lowest resolution axis, for anisotropic cases), a target, downsampled resolution between two (or whichever resolution the most anisotropic axis boasts) and seven mm is selected. The downsampling process is ultimately carried out via resampling to a new downsampled grid, followed by Gaussian smoothing, to mimic real slice spacing and thickness.

As a suitable baseline, we select UniRes [2], a registration-based multi-modal SR technique. UniRes offers pre-trained models on its open-source GitHub repos-

**Table 3.** Mean PSNR for the SR experiment, for all combinations of MRI modalities, including at least the reconstructed modality on the SABRE dataset. •: Modality present, ○: Modality absent. Bold values represent statistically best performances via signed-rank Wilcoxon test.

| Regenerated modality | Input modalities | | | PSNR | |
|---|---|---|---|---|---|
| | FLAIR | MPRAGE | T2-SE | UniRes | Ours |
| FLAIR | • | ○ | ○ | 24.630 (2.773) | **29.561 (2.854)** |
| | • | • | ○ | 25.217 (3.055) | **28.980 (2.614)** |
| | • | ○ | • | 25.326 (2.920) | **29.201 (2.672)** |
| | • | • | • | 24.883 (2.785) | **28.809 (2.748)** |
| MPRAGE | ○ | • | ○ | 25.778 (2.980) | **31.742 (2.324)** |
| | • | • | ○ | 27.381 (3.528) | **31.856 (2.057)** |
| | ○ | • | • | 27.361 (3.780) | **31.624 (2.164)** |
| | • | • | • | 26.947 (3.213) | **32.129 (2.206)** |
| T2-SE | ○ | ○ | • | 25.986 (2.597) | **31.407 (2.430)** |
| | • | ○ | • | 27.217 (3.278) | **31.927 (2.119)** |
| | ○ | • | • | 26.443 (3.011) | **31.695 (2.041)** |
| | • | • | • | 26.676 (3.126) | **32.318 (1.733)** |

itory, which we use for our analyses. Since both methods are multimodal, we test every combination of input modalities for both validations.

Figure 2 shows an example set of high-resolution same-subject multi-modality images, their downsampled counterparts, and the SR-MPMGen and baselines super-resolved outputs. SR-MPMGen is adept at leveraging complementary high-resolution information across modalities, resulting in a reconstruction that does not exhibit the oft-present hallucinatory signs of most SR techniques. Image quality is quantitatively assessed using the peak signal-to-noise ratio (PSNR), a commonplace image reconstruction metric. Results are found in Table PSNR. Across all modalities and input combinations, SR-MPMGen significantly outperforms UniRes, backing up the aforementioned qualitative findings.

**Assessing domain generalisation via a classification downstream:** The work is largely motivated by the fact that the models we train are susceptible to domain shifts and that being able to transition to invariant, quantitative data would alleviate this issue by being domain generalisable. Note the important distinction between domain adaptation and domain generalisation. Domain adaptation techniques *adapt* to the target domain, necessitating target domain data by definition. Conversely, a domain-generalisation technique should be innately compatible with unseen data, making it more versatile and not requiring further processing or training using domain samples.

We illustrate this point via a classification experiment. We train a 2D model using FLAIR images from the BraTS dataset (denoted Qual-FLAIR) to identify the presence of a tumour (>10 pixels) and assess performance on different test modalities. In tandem, we train a model using MPMs derived from only FLAIR images from the BraTS dataset (denoted MPM-FLAIR) and test the model on MPMs derived using different individual modalities, likewise assessing perfor-

**Table 4.** Classification results for tumour detection for Qual-FLAIR and Qual-MPM, evaluated on FLAIR (IoD), MPRAGE (OoD), and SE (OoD) on the BraTs dataset.

| Input modalities | ROC-AUC | |
|---|---|---|
| | Qual-FLAIR | MPM-FLAIR |
| FLAIR | 0.933 | 0.928 |
| SE | 0.723 | **0.786** |
| MPRAGE | 0.692 | **0.753** |

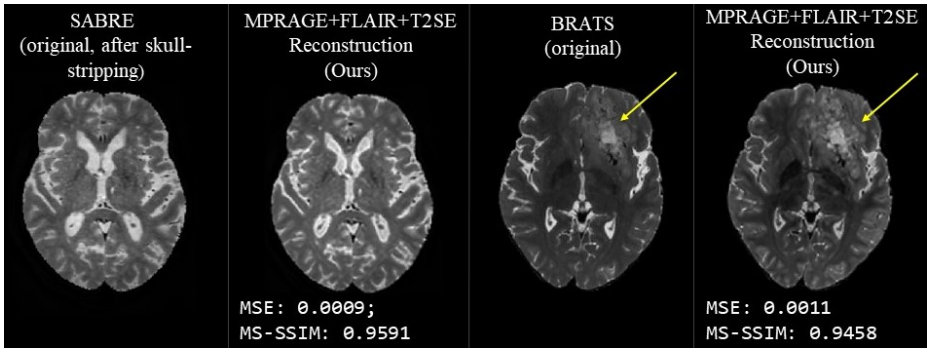


**Fig. 3.** Qualitative showcase of SR-MPMGen reconstructions on OoD skull-stripped data. Arrows denote tumours.

mance. The latter model should be less susceptible to modality sequestration, as it is meant to be invariant. Our architecture is a DenseNet (for details, see supplementary materials). The results of this experiment are shown in Table 4.

Statistical significance between ROC-AUCs of the two models is ascertained via DeLong tests [5]. While both models have diminished performance on MPRAGE and SE validations, MPM-FLAIR is statistically more robust to domain shift, supporting our conjecture that translating images into the isotropic MPM domain leads to stronger domain generalisation. Note that the model has been applied on 1) severely pathological (the model was trained on largely tumour-absent data), 2) skull-stripped (the model was trained on non-skull-stripped data), and 3) sequence-parameter absent data (knowledge of parameters constrains the MPM-translation process, incorrect values will incorrectly skew values, and these had to be estimated for BraTS); the proposed results therefore likely represent a performance floor. Figure 3 compares skull-stripped SABRE T2SE and skull-stripped BraTS T2SE reconstructions using our model. While there are some clear contrast mismatches for BraTS, the MSE and MS-SSIM are not dissimilar from SABRE, affirming the model's versatility.

## 4 Discussion and Conclusion

We present a method for translating various standard MR images into an isotropic, quantitative domain that trains fully self-supervised. We show that adding super-

resolution to the original quantitative translation method does not impair its original functionality but only enhances it. Our downstream classification task evinces our principal motivation is justifiable, namely that versatility is gained from training models using data translated into an invariant domain rather than relying on models whose performance is inextricably tied to its training data. While the current formulation is not compatible with other modalities, it lends itself to expansion by adding further modality-specific branches, something we will seek to explore.

## References


1. Borges, P., Fernandez, V., Tudosiu, P.D., Nachev, P., Ourselin, S., Cardoso, M.J.: Unsupervised Heteromodal Physics-Informed Representation of MRI Data: Tackling Data Harmonisation, Imputation and Domain Shift pp. 53–63 (2023). https://doi.org/10.1007/978-3-031-44689-4_6, https://dl.acm.org/doi/10.1007/978-3-031-44689-4_6
2. Brudfors, M., Balbastre, Y., Nachev, P., Ashburner, J.: Mri super-resolution using multi-channel total variation. In: Annual Conference on Medical Image Understanding and Analysis. pp. 217–228. Springer (2018)
3. Cardoso, M.J., Modat, M., Vercauteren, T., Ourselin, S.: Scale Factor Point Spread Function Matching: Beyond Aliasing in Image Resampling, p. 675–683. Springer International Publishing (2015). https://doi.org/10.1007/978-3-319-24571-3_81, http://dx.doi.org/10.1007/978-3-319-24571-3_81
4. Deichmann, R., Good, C.D., Josephs, O., Ashburner, J., Turner, R.: Optimization of 3-D MP-RAGE sequences for structural brain imaging. NeuroImage **12**(1), 112–127 (2000). https://doi.org/10.1006/NIMG.2000.0601, https://pubmed.ncbi.nlm.nih.gov/10875908/
5. Delong, E.R., Delong, D.M., Clarke-Pearson, D.L.: Comparing the Areas under Two or More Correlated Receiver Operating Characteristic Curves: A Nonparametric Approach **44**(3), 837–845 (1988), https://about.jstor.org/terms
6. Hornak, J.: The Basics of MRI, https://www.cis.rit.edu/htbooks/mri/
7. Iglesias, J.E., Billot, B., Balbastre, Y., Magdamo, C., Arnold, S.E., Das, S., Edlow, B.L., Alexander, D.C., Golland, P., Fischl, B.: SynthSR: A public AI tool to turn heterogeneous clinical brain scans into high-resolution T1-weighted images for 3D morphometry. Science Advances **9**(5) (feb 2023). https://doi.org/10.1126/sciadv.add3607, https://www.science.org/doi/10.1126/sciadv.add3607
8. Jacobs, L., Mandija, S., Liu, H., van den Berg, C.A., Sbrizzi, A., Maspero, M.: Generalizable synthetic MRI with physics-informed convolutional networks. Medical physics (2023). https://doi.org/10.1002/MP.16884, https://pubmed.ncbi.nlm.nih.gov/38063208/
9. Jia, X., Liu, Y., Yang, Z., Yang, D.: Multi-modality self-attention aware deep network for 3D biomedical segmentation. BMC Medical Informatics and Decision Making **20**(3), 1–7 (jul 2020). https://doi.org/10.1186/s12911-020-1109-0
10. Jog, A., et al.: Mr image synthesis by contrast learning on neighborhood ensembles. Medical image analysis **24**(1), 63–76 (2015)
11. Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber,

M.A., Arbel, T., Avants, B.B., Ayache, N., Buendia, P., Collins, D.L., Cordier, N., Corso, J.J., Criminisi, A., Das, T., Delingette, H., Demiralp, Ç., Durst, C.R., Dojat, M., Doyle, S., Festa, J., Forbes, F., Geremia, E., Glocker, B., Golland, P., Guo, X., Hamamci, A., Iftekharuddin, K.M., Jena, R., John, N.M., Konukoglu, E., Lashkari, D., Mariz, J.A., Meier, R., Pereira, S., Precup, D., Price, S.J., Raviv, T.R., Reza, S.M., Ryan, M., Sarikaya, D., Schwartz, L., Shin, H.C., Shotton, J., Silva, C.A., Sousa, N., Subbanna, N.K., Szekely, G., Taylor, T.J., Thomas, O.M., Tustison, N.J., Unal, G., Vasseur, F., Wintermark, M., Ye, D.H., Zhao, L., Zhao, B., Zikic, D., Prastawa, M., Reyes, M., Van Leemput, K.: The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). IEEE transactions on medical imaging **34**(10), 1993–2024 (oct 2015). https://doi.org/10.1109/TMI.2014.2377694, https://pubmed.ncbi.nlm.nih.gov/25494501/

12. Modat, M., Cash, D.M., Daga, P., Winston, G.P., Duncan, J.S., Ourselin, S.: Global image registration using a symmetric block-matching approach. Journal of medical imaging (Bellingham, Wash.) **1**(2), 024003 (sep 2014). https://doi.org/10.1117/1.JMI.1.2.024003, https://pubmed.ncbi.nlm.nih.gov/26158035/
13. Plenge, E., Poot, D.H., Bernsen, M., Kotek, G., Houston, G., Wielopolski, P., Van Der Weerd, L., Niessen, W.J., Meijering, E.: Super-resolution methods in MRI: Can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time? Magnetic Resonance in Medicine **68**(6), 1983–1993 (dec 2012). https://doi.org/10.1002/MRM.24187, https://onlinelibrary.wiley.com/doi/full/10.1002/mrm.24187 https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24187 https://onlinelibrary.wiley.com/doi/10.1002/mrm.24187
14. Tillin, T., et al.: Southall And Brent REvisited: Cohort profile of SABRE, a UK population-based comparison of cardiovascular disease and diabetes in people of European, Indian Asian and African Caribbean origins. International journal of epidemiology **41**(1), 33–42 (feb 2012). https://doi.org/10.1093/IJE/DYQ175, https://pubmed.ncbi.nlm.nih.gov/21044979/
15. Torralba, A., Efros, A.A.: Unbiased look at dataset bias. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition pp. 1521–1528 (2011). https://doi.org/10.1109/CVPR.2011.5995347
16. Varadarajan, D., Bouman, K.L., van der Kouwe, A., Fischl, B., Dalca, A.V.: Unsupervised learning of MRI tissue properties using MRI physics models (jul 2021). https://doi.org/10.48550/arxiv.2107.02704, https://arxiv.org/abs/2107.02704v1