

MMORF—FSL’s MultiMODal Registration Framework

Frederik J. Lange^{a,*}, Christoph Arthofer^a, Andreas Bartsch^b, Gwenaëlle Douaud^a, Paul McCarthy^a, Stephen M. Smith^a, Jesper L. R. Andersson^a

^a *Wellcome Centre for Integrative Neuroimaging, FMRIB, Nuffield Department of Clinical Neurosciences, University of Oxford, United Kingdom*

^b *Department of Neuroradiology, University of Heidelberg, Germany*

Abstract

We present MMORF—FSL’s MultiMODal Registration Framework—a newly released nonlinear image registration tool designed primarily for application to MRI images of the brain. MMORF is capable of simultaneously optimising both displacement and rotational transformations within a single registration framework by leveraging rich information from multiple scalar and tensor modalities. The regularisation employed in MMORF promotes local rigidity in the deformation, and we have previously demonstrated how this effectively controls both shape and size distortion, and leads to more biologically plausible warps. The performance of MMORF is benchmarked against three established nonlinear registration methods—FNIRT, ANTs and DR-TAMAS—across four domains: FreeSurfer label overlap, DTI similarity, task-fMRI cluster mass, and distortion. Results show that MMORF performs as well as or better than all other methods across every domain—both in terms of accuracy and levels of distortion. MMORF is available as part of FSL, and its inputs and outputs are fully compatible with existing workflows. We believe that MMORF will be a valuable tool for the neuroimaging community, regardless of the domain of any downstream analysis, providing state-of-the-art registration performance that integrates into the rich and widely adopted suite of analysis tools in FSL.

Keywords: registration, nonlinear, multimodal, volumetric, FSL

^{*}Corresponding Author

Email address: frederik.lange@ndcn.ox.ac.uk (Frederik J. Lange)

1. Introduction

2 In this paper we describe and evaluate FSL’s MultiMODal Registration
3 Framework (MMORF). MMORF is a nonlinear image registration tool, pri-
4 marily intended for magnetic resonance imaging (MRI) of the brain.

5 Biomedical image registration is a core component in most neuroimaging
6 processing pipelines. If we assume that all brains, regardless of appear-
7 ance, are built using the same anatomical components arranged in the same
8 configuration (*i.e.*, they are topologically identical), then we can use image
9 registration to find the set of deformations that map all brains to a com-
10 mon reference space/template brain in a one-to-one manner. In this case, a
11 template may refer to either a group-average brain or an individual subject
12 brain. Although this assumption may not always hold, it is valid enough
13 for registration to be used in both the localisation and quantification of sim-
14 ilarities and differences between individuals or population groups—that is,
15 the template allows us to say *where* the differences occur, and the mappings
16 ensure that *what* we are comparing is measured at the same place in all sub-
17 jects. A poor registration will, therefore, impact both the power to detect
18 and the ability to accurately localise any effects of interest across subjects.

19 However, the reality is that image registration is an ill-posed problem,
20 and therefore the true, one-to-one mapping can generally never be found.
21 By ill-posed, we mean that there is typically not enough information in an
22 image itself to find a mapping that uniquely maximises some measure of sim-
23 ilarity between a subject’s brain and the template. Therefore, registration
24 methods require regularisation to constrain the solution to be unique. Reg-
25 ularisation achieves this by encoding a model of which types of deformations
26 are considered more likely than others. The challenge in designing a reg-
27 istration tool is then to find the best way to combine *image information* and
28 *regularisation* to produce as good an approximation to the true mapping as
29 possible.

30 MMORF addresses the image information aspect of this challenge by
31 taking a multimodal approach to computing brain similarity. Up until now we
32 have described registration as if there is only one measurement/image of the
33 brain that we can use to find the correct mapping. However, with MRI we are
34 able to acquire a number of different image modalities—all within the same
35 imaging session, and each with different contrast and information content.
36 In MMORF, we have leveraged this complementary information to reduce
37 the degree to which the registration problem is ill-posed, thereby improving

38 our confidence that we are finding an accurate mapping for each subject.
39 MMORF’s registration abilities extend beyond scalar modalities to include
40 diffusion tensor imaging (DTI). When using the full tensor data (rather than
41 scalar, rotationally-invariant, derived features such as fractional anisotropy
42 (FA)), it matches the directional information in the diffusion tensor to guide
43 the alignment of, in particular, white matter.

44 MMORF’s regularisation method is one of its most unique attributes. By
45 employing a penalty that aims to preserve the original shape and volume of
46 the data as far as is reasonable, it is able to produce deformations that are far
47 more biologically plausible than those generated by conventional techniques.
48 A detailed description of its implementation in MMORF, and a thorough
49 evaluation of the benefits of this form of regularisation, can be found in
50 Lange et al. (2020).

51 The cost of combining multiple modalities with a complex regularisation
52 model is that the computational requirements of the method increase. We
53 have addressed this from the outset by designing MMORF to use graphics
54 processing units (GPUs) to parallelise its execution. This allows MMORF to
55 execute 1 mm isotropic registrations with reasonable runtimes of between 5
56 and 45 minutes on modern hardware, depending on the number of modalities
57 used.

58 Through reducing the degree to which the registration problem is ill-
59 posed (using multimodal data), and producing more realistic deformations
60 in those regions where it is (using biomechanically realistic regularisation),
61 MMORF is capable of state-of-the-art registration accuracy. In the methods
62 section below we detail the most important design decisions made when de-
63 veloping MMORF, so that how it operates can be clearly understood. We
64 then contextualise these decisions with reference to a set of comparable cur-
65 rent registration tools. Finally, we validate and benchmark MMORF against
66 these tools in order to demonstrate its performance and utility.

67 2. Methods

68 As stated in the introduction, the true one-to-one mapping between brain
69 images can never be known, and therefore image registration can only find the
70 “optimal” mapping based on the available information and our prior beliefs
71 about what mappings are more likely than others. Unsurprisingly then,
72 registration is normally formulated as an optimisation problem, requiring a
73 cost function to minimise. MMORF is intended to be used with any number

74 of scalar and tensor images driving the registration, and we therefore choose
75 to minimise a total cost function \mathcal{C}_{tot} that is defined as follows:

$$\mathcal{C}_{tot}(\mathbf{w}) = \sum_{\mathbf{x} \in \Omega} \left(\sum_s \lambda_s \mathcal{C}_s(\mathbf{w}, \mathbf{x}) + \sum_t \lambda_t \mathcal{C}_t(\mathbf{w}, \mathbf{x}) + \lambda_r \mathcal{C}_r(\mathbf{w}, \mathbf{x}) \right) \quad (1)$$

76 where \mathbf{w} are the warp parameters being optimised, Ω is the domain over
77 which the warp is defined, λ_* are cost function weightings, and subscripts s ,
78 t , and r refer to scalar, tensor and regularisation respectively.

79 Each scalar and tensor cost function is based on an image dissimilarity
80 metric between a pair of images—one belonging to the reference subject (of
81 ten a template), and one to the moving subject. Although Equation 1 shows
82 that these cost functions are separable, and can therefore be evaluated in-
83 dependently, it is critical that *within* each subject all images are accurately
84 co-registered. If not, then a single warp cannot correctly map all modalities
85 between subjects, and registration accuracy will suffer. We do not attempt
86 to ensure this within MMORF itself. Rather, we rely on running accurate,
87 rigid, *between*-modality, *within*-subject registration using FSL FLIRT (Jen-
88 kinson and Smith, 2001; Jenkinson et al., 2002), as well as distortion correction
89 for relevant modalities such as DTI (Andersson et al., 2003; Andersson and
90 Sotiroopoulos, 2016), before performing nonlinear registration. However, it is
91 not necessary to resample any of the images following rigid alignment be-
92 fore feeding them into MMORF. Instead, MMORF expects a separate affine
93 transformation matrix to be supplied for each image being registered (in
94 both the moving and reference subject datasets). This affine transformation
95 points to a separate, user-specified, “warp-space image”, whose extents de-
96 fine the domain over which MMORF will estimate and output its nonlinear
97 deformation. This eliminates the need for multiple resampling of the data,
98 and requires no matching of the image resolution or dimensions of any of the
99 images being registered.

100 With this information in hand, the following sections elaborate on some
101 of the key decisions that went into MMORF’s design, and the role they play
102 in its performance.

103 2.1. Optimisation Strategy

104 Nonlinear image registration tools generally use one of two iterative op-
105 timisation approaches—first or second order. First order methods consider

106 only the gradient of the cost function, with respect to each optimisable pa-
 107 rameter, when picking a parameter update step. Second order methods ex-
 108 tend this by also considering how long the derivative is valid for and the
 109 interaction between parameters. First order algorithms tend to have up-
 110 date steps that can be calculated more quickly, but second order algorithms
 111 tend to require fewer update steps to reach convergence. In our experience,
 112 the trade-off tends to favour second order approaches for methods such as
 113 MMORF.

114 We have, therefore, implemented two variants of the (second-order) Gauss-
 115 Newton (GN) optimisation strategy—which is itself a variation on Newton’s
 116 method.

Newton’s method is an iterative optimisation algorithm that uses a quadratic Taylor approximation of the cost function \mathcal{C} around the current set of parameters \mathbf{w} (Nocedal and Wright, 2006, Ch.10, pp 254). That is:

$$\mathcal{C}(\mathbf{w} + \Delta\mathbf{w}) \approx \mathcal{C}(\mathbf{w}) + \nabla\mathcal{C}(\mathbf{w})^\top \Delta\mathbf{w} + \frac{1}{2} \Delta\mathbf{w}^\top \mathbf{H}(\mathbf{w}) \Delta\mathbf{w} \quad (2)$$

where $\nabla\mathcal{C}$ and \mathbf{H} are the gradient and the Hessian of the cost function, respectively. The update $\Delta\mathbf{w}$ that minimises this approximation to the cost function (when the cost function is convex) is then:

$$\Delta\mathbf{w} = -\mathbf{H}^{-1}(\mathbf{w}) \nabla\mathcal{C}(\mathbf{w}) \quad (3)$$

For cost functions that can be written in the form:

$$\mathcal{C}(\mathbf{w}) = \frac{1}{2} \sum_{i=1}^N c_i^2(\mathbf{w}) \quad (4)$$

The gradient is then:

$$\nabla\mathcal{C}(\mathbf{w}) = \sum_{i=1}^N c_i(\mathbf{w}) \frac{\partial c_i}{\partial \mathbf{w}} \quad (5)$$

And the Hessian is:

$$\mathbf{H}(\mathbf{w}) = \sum_{i=1}^N \frac{\partial c_i}{\partial \mathbf{w}}^\top \frac{\partial c_i}{\partial \mathbf{w}} + c_i(\mathbf{w}) \frac{\partial^2 c_i}{\partial \mathbf{w}^2} \quad (6)$$

117 Gauss-Newton approximates \mathbf{H} by dropping the second, mixed partial derivative,
 118 term in Equation 6. Close to the optimum set of \mathbf{w} , the Newton and
 119 Gauss-Newton Hessians tend towards equivalence. Further from the optimum,
 120 the Gauss-Newton Hessian has the benefit of being symmetric positive
 121 definite, and therefore the approximation of the cost function is always
 122 convex. This ensures that each step is always in a direction that would improve
 123 the original cost function.

However, the Gauss-Newton step can still lead to an increase in the cost function if it oversteps. This can be addressed by the Levenberg-Marquardt (LM) extension (Nocedal and Wright, 2006, Ch.10, pp258), which replaces the Gauss-Newton Hessian (\mathbf{H}_{GN}) with:

$$\mathbf{H}_{LM} = \mathbf{H}_{GN} + \lambda \mathbf{I}. \quad (7)$$

124 For small values of λ , the update $\Delta\mathbf{w}$ behaves much like a Gauss-Newton
 125 update, and for large values of λ it behaves more like gradient descent. After
 126 choosing a sensible starting value for λ (e.g., 0.001), a typical Levenberg-
 127 Marquardt iteration then proceeds as follows:

- 128 • Calculate a candidate parameter update $\Delta\mathbf{w}$ using Equation 3.
- 129 • If $\mathcal{C}(\mathbf{w} + \Delta\mathbf{w}) < \mathcal{C}(\mathbf{w})$ accept the update and set $\lambda = \lambda \div 10$.
- 130 • If $\mathcal{C}(\mathbf{w} + \Delta\mathbf{w}) > \mathcal{C}(\mathbf{w})$ set $\lambda = \lambda \times 10$ and recompute $\Delta\mathbf{w}$ and $\mathcal{C}(\mathbf{w} + \Delta\mathbf{w})$,
 131 repeating until $\mathcal{C}(\mathbf{w} + \Delta\mathbf{w}) < \mathcal{C}(\mathbf{w})$ and then accept the update.

132 The Majorise-Minimisation (MM) algorithm (Hunter and Lange, 2004)
 133 is a method for cost function minimisation that can be used when storing
 134 the full Gauss-Newton Hessian is infeasible (e.g., due to memory constraints,
 135 which prove to be important in this application). In essence, what this
 136 algorithm states is that if $z(\mathbf{w}) \geq y(\mathbf{w}), \forall \mathbf{w}$, with both y and z convex
 137 and touching, then $\mathbf{w}_{k+1} = \arg \min_{\mathbf{w}} z(\mathbf{w})$ will either reduce y or leave it
 138 unchanged.

Since the quadratic Taylor approximation of $\mathcal{C}(\mathbf{w})$ using \mathbf{H}_{GN} in Equation 2 is convex, it may serve as y . Chun and Fessler (2018) showed that because the diagonal elements of \mathbf{H}_{GN} are all positive:

$$\mathbf{H}_{MM} = \text{diag}(|\mathbf{H}_{GN}| \mathbf{1}) \succeq \mathbf{H}_{GN} \quad (8)$$

139 where $|\mathbf{H}_{GN}|$ is the matrix of the absolute values of \mathbf{H}_{GN} and $\mathbf{1}$ is column
 140 vector of ones. The \succeq symbol means that \mathbf{H}_{MM} is at least as positive definite

141 as \mathbf{H}_{GN} . Therefore, substituting \mathbf{H}_{MM} for \mathbf{H}_{GN} in the quadratic Taylor
142 approximation of $\mathcal{C}(\mathbf{w})$ majorises y , and may serve as z . In other words
143 simply, replacing \mathbf{H}_{GN} with \mathbf{H}_{MM} in the LM algorithm and solving as usual
144 satisfies the MM algorithm requirements.

145 The major advantage of using \mathbf{H}_{MM} is that it is only non-zero along its
146 main diagonal, and therefore requires far less memory to store than \mathbf{H}_{GN} .
147 This is particularly important when the number of parameters in \mathbf{w} is large,
148 as it may not be possible to store \mathbf{H}_{GN} .

149 The reason for including both the Levenberg-Marquardt and Majorise-
150 Minimisation modifications of Gauss-Newton, is that MMORF both calcu-
151 lates the Hessian and solves for the update step on GPU hardware. GPUs
152 typically have a limited amount of RAM, and, therefore, storing the full
153 Gauss-Newton Hessian becomes impossible for warps beyond a certain res-
154 olution. At this point, MMORF switches to Majorise-Minimisation, which
155 requires on the order of 1000 times less memory. In practice, Majorise-
156 Minimisation requires more steps to converge, but requires no other changes
157 to the registration algorithm, making it an appealing option.

158 *2.2. Transformation Model*

159 MMORF employs a transformation of the small deformation framework
160 variety (Bajcsy and Kovačič, 1989; Miller et al., 1997)—that is, it defines the
161 deformation as a displacement field, rather than a velocity field as per the
162 large deformation framework (Bro-Nielsen and Gramkow, 1996; Miller et al.,
163 1997). A major difference between these two frameworks is their relationship
164 to diffeomorphism.

165 Diffeomorphism is a desirable property in image registration. Diffeomor-
166 phic transformations are smooth, one-to-one, and onto, and are therefore
167 guaranteed to induce neither folding nor tearing when applied to an image.
168 The large deformation framework, and in particular the large deformation
169 diffeomorphic metric mapping (LDDMM) (Beg et al., 2005) family of tools,
170 have the advantage that their transformation model can be made inherently
171 diffeomorphic—i.e., they are, in principle, diffeomorphic by construction.

172 However, despite not being inherently diffeomorphic, the small deforma-
173 tion framework can be made diffeomorphic by employing a regularisation
174 penalty that enforces diffeomorphism (Rohlfing et al., 2003). This has the
175 benefit that warp-induced distortions, such as changes in shape and volume,
176 can be calculated (and therefore controlled) directly from the model param-
177 eters themselves. In contrast, large deformation models require the vector

178 field that they parametrise to first be integrated over time and converted into
179 a displacement field before being able to calculate such measures. Therefore,
180 explicitly *controlling* the amount of distortion is harder in the large deforma-
181 tion framework, despite its guarantee of diffeomorphism. This is important in
182 our case, as the advantage of MMORF’s regularisation (described in Section
183 2.4) is that it nonlinearly penalises the stretching and compressing effects of
184 the warp directly.

185 MMORF’s transformation is parametrised using cubic B-splines (Unser
186 et al., 1993a,b). This imposes an inherent smoothness to the deformations,
187 as cubic B-splines, and therefore the warps, have C2 continuity. The warp
188 field is then also well defined across all of the image (*i.e.*, not just at voxel
189 centres), eliminating any need to match image subsampling or resolution ei-
190 ther before or during registration. Despite their smoothness, B-splines still
191 have compact support (*i.e.*, the effect of a particular spline is exactly zero
192 at a fixed distance from the spline centre), and therefore interaction effects
193 between splines are fixed and finite. This has important implications for
194 second order optimisation (such as the Gauss-Newton method employed by
195 MMORF), since it means that the Hessian matrix (which encodes the inter-
196 action between optimisable parameters) is sparse and predictably patterned.

197 Furthermore, the spatial derivatives of B-splines are smooth and have a
198 closed form solution. Since these derivatives are required to calculate gradi-
199 ents and Hessians for the optimisation of warps in MMORF, this is compu-
200 tationally beneficial.

201 We now provide a more explicit description of how we have formulated our
202 transformation model, and how this interacts with calculation of the Hessian
203 during optimisation (the most computationally intensive part of the registra-
204 tion algorithm). A set of B-spline basis functions can be used to transform
205 a set of sample coordinates in a reference image f to their corresponding
206 coordinates in a moving image g as follows:

f = Reference image

g = Moving image

N = Number of sampled voxels in f

x/y/z = x , y and z coordinates of samples in f

$$\mathbf{X} = \underbrace{\begin{bmatrix} \mathbf{x}^\top \\ \mathbf{y}^\top \\ \mathbf{z}^\top \end{bmatrix}}_{3 \times N} = \begin{bmatrix} x_0 & x_1 & \cdots & x_N \\ y_0 & y_1 & \cdots & y_N \\ z_0 & z_1 & \cdots & z_N \end{bmatrix}$$

M = Number of splines per warp direction

$\mathbf{w}_x/\mathbf{w}_y/\mathbf{w}_z = x, y$ and z direction warp parameters

$$\mathbf{W} = \underbrace{\begin{bmatrix} \mathbf{w}_x^\top \\ \mathbf{w}_y^\top \\ \mathbf{w}_z^\top \end{bmatrix}}_{3 \times M} = \begin{bmatrix} w_{x0} & w_{x1} & \cdots & w_{xM} \\ w_{y0} & w_{y1} & \cdots & w_{yM} \\ w_{z0} & w_{z1} & \cdots & w_{zM} \end{bmatrix}$$

\mathbf{b}_m = Vectorised m^{th} B-spline at sample positions \mathbf{X}

$$\mathbf{B} = \underbrace{\begin{bmatrix} \mathbf{b}_0^\top \\ \mathbf{b}_1^\top \\ \vdots \\ \mathbf{b}_M^\top \end{bmatrix}}_{M \times N} = \underbrace{\begin{bmatrix} b_{00} & b_{01} & \cdots & b_{0N} \\ b_{10} & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ b_{M0} & \cdots & \cdots & b_{MN} \end{bmatrix}}_{1 \text{ voxel per column}} \quad \left. \begin{array}{l} \text{1 spline per row} \\ \text{1 voxel per column} \end{array} \right\}$$

$\phi(\mathbf{X}, \mathbf{W})$ = Transformed sample coordinates

$$= \mathbf{X} + \mathbf{WB}$$

207 Note that here we assume that the warp space and the reference image
 208 f are the same, and that f and g are already affinely registered to each
 209 other. We use the same configuration of B-splines to define the warps in all
 210 3 directions (*i.e.*, the number, order and spatial extent of splines defining
 211 the displacements in each direction (x , y and z) is the same, and they are
 212 located identically in space). M is chosen to be the set of splines whose
 213 spatial support, wholly or partially, overlaps with the domain over which f
 214 is defined. The warp parameters $\mathbf{w}_{x/y/z}$ are then the coefficients of each B-
 215 spline, and each parameter only affects the displacement in a single direction.

216 The compact support of B-splines means that \mathbf{B} is very sparse. Additionally,
 217 (disregarding edge cases) each row of \mathbf{B} is simply a shifted version
 218 of any other row. We therefore never store \mathbf{B} explicitly, and instead compute
 219 \mathbf{WB} using convolution.

Another benefit of the sparsity in \mathbf{B} is that it induces sparsity in the Hessian of cost functions based on this transformation. For a mean squared error cost function, each element of the Hessian can be calculated as:

$$\mathbf{H}_{ij} = \sum_{\substack{\mathbf{x} \in \mathbf{b}_i \cap \mathbf{b}_j \\ \text{spline intersection}}} \mathbf{b}_i(\mathbf{x}) \mathbf{b}_j(\mathbf{x}) \underbrace{\frac{\partial g}{\partial xyz(i)}(\phi(\mathbf{x})) \frac{\partial g}{\partial xyz(j)}(\phi(\mathbf{x}))}_{\text{gradient image}}, \quad (9)$$

220 where $\frac{\partial g}{\partial xyz(i)}$ means “differentiate with respect to the direction in which B-
 221 spline i causes displacement to occur”. As each entry in the Hessian repre-
 222 sents the interaction between two of the B-spline basis functions, only those
 223 combinations of splines that overlap in their spatial support will ever produce
 224 non-zero entries. The number of non-zero entries per row of the Hessian is at
 225 most $3 \times 7^3 = 1026$ for 3D images and cubic B-splines. Considering that the
 226 number of parameters being optimised over can easily exceed 10^6 , this would
 227 lead to a matrix that is at least 99.9% sparse. The redundancy increases
 228 as the warp resolution increases (the knot-spacing is reduced), which means
 229 that the more parameters one attempts to estimate, the sparser the Hessian
 230 becomes. Therefore, the memory requirement for the Hessian scales much
 231 better with resolution than might initially be feared.

232 *2.3. Image Cost Functions*

233 As stated earlier, MMORF optimises the cost function defined in Equa-
 234 tion 4, which is the sum over individual cost functions for each pair of images.
 235 In this section we describe the choice of cost function for scalar and tensor im-
 236 ages, as well as how MMORF implements cost function masking/weighting.

237 *2.3.1. Scalar*

Scalar image dissimilarity is calculated using the mean squared error (MSE) across the image. That is, for two scalar images f and g :

$$\mathcal{C}_s(\mathbf{w}, f, g) = \frac{1}{\Omega} \sum_{\mathbf{x} \in \Omega} (f(\mathbf{x}) - g(\mathbf{x}, \mathbf{w}))^2 \quad (10)$$

238 MSE is ideally suited to the GN family of optimisation methods, which have
239 an implicit assumption that the cost function being optimised is some form
240 of squared difference.

241 Robust mean intensity estimation is used to scale each image separately in
242 order to account for linear scaling intensity differences between image pairs.
243 When using MSE cost functions, spatial intensity inhomogeneities (*e.g.*, due
244 to transmit and receive bias fields in MRI (Vovk et al., 2007; Andersson et al.,
245 2019)) can be mistakenly interpreted as image misalignment, leading to un-
246 necessary (and incorrect) image warping. In MMORF we therefore provide
247 the option to explicitly model such inhomogeneities as a smoothly varying
248 multiplicative field acting on the reference image. This can be enabled or dis-
249 abled for each image pair independently. The bias field is parametrised using
250 cubic B-splines, and the resolution (knot-spacing) and smoothness (bending
251 energy (Bookstein, 1997)) are set on a per image pair basis.

252 *2.3.2. Tensor*

253 Tensor image dissimilarity is calculated using the mean squared Frobenius
254 norm (MSFN) across the image. This is exactly equivalent to summing the
255 MSE for each of the 9 elements of the diffusion tensor, and therefore fits our
256 GN optimisation strategy. That is, for two tensor images \mathbf{F} and \mathbf{G} :

$$\mathcal{C}_t(\mathbf{w}, \mathbf{F}, \mathbf{G}) = \frac{1}{\Omega} \sum_{\mathbf{x} \in \Omega} \sum_{i=1}^3 \sum_{j=1}^3 (\mathbf{F}_{ij}(\mathbf{x}, \mathbf{w}) - \mathbf{G}_{ij}(\mathbf{x}, \mathbf{w}))^2 \quad (11)$$

257 Note that, in contrast to the scalar case, the values in the reference image
258 \mathbf{F} are also a function of the warp parameters \mathbf{w} . This is because MMORF
259 applies the *rotational* effect of the warp only to the reference image, as well
260 as the usual *displacement* effect only on the moving image.

261 Rotation of tensors uses the finite strain (FS) method (Alexander et al.,
262 2001). The Jacobian matrix of the warp at each position $\mathbf{J}(\mathbf{x}, \mathbf{w})$ represents
263 the first order linear approximation of the deformation at that point, and
264 from it the local rotational effect of the warp $\mathbf{R}(\mathbf{x}, \mathbf{w})$ can be calculated
265 using:

$$\mathbf{R} = (\mathbf{J}\mathbf{J}^\top)^{-\frac{1}{2}}\mathbf{J} \quad (12)$$

266 We include the effects of both displacement and rotation when calculating
267 the gradient and the Hessian of the tensor cost function, using the derivation
268 of Yeo et al. (2009) for the closed-form of the derivative due to rotation.

269 Since the diffusion tensor model is quantitative, no intensity rescaling is
270 performed. This implicitly assumes that the tensors are always represented
271 with the same units. Similarly, fitting of the diffusion tensor compensates for
272 any intensity inhomogeneity present in the raw diffusion-weighted images,
273 and therefore no bias field estimation is required .

274 MMORF assumes that diffusion tensors are stored in FSL dtifit¹ for-
275 mat—*i.e.*, a 4D volume where the 4th dimension contains the 6 upper-diagonal
276 elements vectorised row-wise, and the x-direction is defined in radiological
277 convention (R-L).

278 *2.3.3. Masking*

279 In certain instances it may be beneficial to focus a registration algorithm
280 on a particular part of an image or, alternatively, for it to ignore a cer-
281 tain region. To this end, we have implemented masking/weighting within
282 MMORF.

283 Masks can be supplied for all reference and moving images independently,
284 and it is required that these masks be in the same space as their corresponding
285 images. The masks are treated as containing voxelwise multiplicative factors
286 that are applied to the cost function during optimisation. Setting any region
287 of a mask to zero will cause the algorithm to ignore the impact of that region’s
288 similarity between reference and moving images, and the deformation in that
289 region will be determined purely by the regularisation. However, masks do
290 not have to be binary, and so a “soft” mask can be used to favour the
291 alignment of one region (*e.g.*, the brain over the rest of the head), without
292 ignoring it completely.

293 In MMORF, masks can be enabled or disabled for each image at every
294 iteration independently, allowing for maximum flexibility. We have, at times,
295 found it useful to have a mask be applied only during higher resolution iter-
296 ations of the registration.

297 *2.4. Regularisation*

298 Since image registration is an under-constrained problem, regularisation
299 is essential to ensure that the resulting warp fields are biologically plausible.

¹<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FDT/UserGuide#DTIFIT>

300 A plausible warp should, at the very least, be diffeomorphic. Diffeomorphic
301 warps are both one-to-one and onto—that is, each point in the reference
302 space maps to a unique point in the moving space, and each point in the mov-
303 ing space is reachable from the reference space. They require the mapping
304 to be continuous, smooth (*i.e.*, have a continuous derivative), and invert-
305 ible (*i.e.*, have a finite, positive Jacobian determinant everywhere). B-spline
306 parametrised warps are, by definition, smooth and continuous. Therefore,
307 provided the Jacobian determinant remains > 0 everywhere, the warp is
308 diffeomorphic.

309 Diffeomorphism is, however, only one desirable trait in a warp. It guar-
310 antees that an image is never torn or folded, but that is all. Typically,
311 penalising the variation of the Jacobian determinant from a value of 1 is
312 used to encourage or ensure (depending on the choice of penalty) that warps
313 remain diffeomorphic in the small deformation framework. However, penal-
314 ising the Jacobian determinant directly only controls volumetric distortion
315 (*i.e.*, changes in an object’s size), and does not in any way control shape dis-
316 tortion (*i.e.*, changes in an object’s shape). The singular values of the local
317 Jacobian represent stretches/compressions along three orthogonal directions.
318 Therefore, the difference between them indicates the degree of shape distor-
319 tion. Because the Jacobian determinant is the product of the singular values
320 of the Jacobian, one can control changes in both size and shape by controlling
321 only the singular values. Penalising deviations of the singular values from 1,
322 and ensuring that none become negative, leads to diffeomorphic warps with
323 desirably little distortion in both volume and shape.

324 In MMORF, the specific penalty used is the mean (across the image) of
325 the sum of the squared logarithms of the singular values at each voxel, as
326 shown in Equation 13 where s_i is the i^{th} singular value of the local Jacobian
327 matrix \mathbf{J} . This is an adaptation of the penalty first proposed by Ashburner
328 et al. (1999, 2000). Its implementation in MMORF is described (and eval-
329 uated) in detail in Lange et al. (2020), where we demonstrate the positive
330 effect that this form of regularisation has on the biological plausibility of the
331 warps MMORF produces.

$$\mathcal{C}_r(\mathbf{w}) = \frac{1}{\Omega} \sum_{\mathbf{x} \in \Omega} \sum_{i=1}^3 (\log s_i(\mathbf{J}(\mathbf{x})))^2 \quad (13)$$

332 By taking the squared logarithm, this penalty tends to infinity as any
333 singular value tends to either zero or infinity. Additionally, this does not

334 penalize any transformations that are locally rigid (*i.e.*, regions that are
335 only translated and/or rotated). Therefore, MMORF regularisation enforces
336 diffeomorphism and encourages local rigidity thereby controlling both volu-
337 metric and shape distortions. The highly nonlinear form of the regularisation
338 also allows the weighting to be set such that large deformations are allowed
339 when necessary, while still ensuring diffeomorphism, thereby overcoming one
340 of the perceived limitations of the small deformation framework.

341 *2.5. Inverse Consistency*

342 An inverse consistent registration algorithm will produce the inverse of the
343 original warp when the reference and moving images swap roles (Christensen,
344 1999). Since we often consider the choice of reference and moving image to
345 be arbitrary (*e.g.*, such as when registering two individuals to each other),
346 this is a desirable property to have.

347 Some tools, such as ANTs (Avants et al., 2008), are inverse consistent
348 by design (excluding the affine initialisation). This is achieved by registering
349 both moving and reference images to a mid-space. However, this essentially
350 means running two registrations per pair of subjects, which takes double the
351 time to compute.

352 In MMORF we have taken a different approach where, rather than en-
353 forcing a symmetric warp, we symmetrise the cost function being minimised
354 instead. This is achieved by multiplying each cost function (both image sim-
355 ilarity and regularisation) by a weighting term of $1 + |\mathbf{J}|$. Since the cost
356 functions are evaluated on a regularly sampled grid in the space of the refer-
357 ence image, the intuition for this weighting is that it always accounts for the
358 total volume in both images to which that value of the cost function applies.
359 In the continuous case, this can be shown to exactly symmetrise the cost
360 function (Tagare et al., 2009). However, since we are dealing with discretely
361 sampled data, this correction is only approximate in MMORF. Nevertheless,
362 this is a better solution than leaving symmetry unaccounted for, and has
363 the computational benefit that only one warp field need be calculated during
364 registration.

365 *2.6. Multiresolution Pyramid*

366 MMORF employs a coarse-to-fine multiresolution optimisation strategy
367 (Bajcsy and Kovačić, 1989; Szeliski and Coughlan, 1997). This has been
368 shown to be beneficial in avoiding local minima during optimisation, as well
369 as accelerating convergence, across a wide variety of nonlinear registration

370 tools (Zhang et al., 2006; Andersson et al., 2007; Ashburner, 2007; Avants
371 et al., 2008; Modat et al., 2010). In principle, such approaches try to match
372 low-frequency image information first, followed by progressively higher fre-
373 quency information at each subsequent level. Traditionally, users are required
374 to specify a downsampling factor and amount of smoothing for each level of
375 the pyramid. The lower-resolution warp is then defined on the coarser, down-
376 sampled reference image grid. A potential pitfall of this approach is that if
377 insufficient smoothing is applied to the image, then the process of down-
378 sampling introduces aliases in the information being aligned due to violation
379 of the Nyquist criterion. MMORF overcomes this by defining the pyramid
380 according to warp resolution (knot spacing) and image smoothing only. An
381 acceptable level of subsampling is then automatically determined by treating
382 the Gaussian smoothing kernel as a low-pass frequency filter. Therefore, the
383 user can iteratively optimise the registration of different frequency content
384 within the image through applying decreasing amounts of smoothing while
385 keeping the warp resolution the same, without any problems with aliasing.
386 In practice, we restrict the downsampled resolution to be at least as high as,
387 and at most four times higher than, the warp resolution at each level of the
388 pyramid.

389 *2.7. Other Implementation Considerations*

390 MMORF is written in C++ and makes extensive use of GPU parallelisa-
391 tion using Nvidia’s CUDA framework (NVIDIA, 2019). Without the use of
392 GPU parallelisation, certain aspects of the registration would be too computa-
393 tionally burdensome to allow MMORF registrations to complete within a
394 reasonable runtime. With GPU acceleration, a typical 1 mm isotropic regis-
395 tration with MMORF takes \approx 10 min for a single scalar, unimodal image pair
396 and \approx 45 minutes for a single scalar, single tensor, multimodal image pair.
397 However, this reliance on CUDA means that MMORF is only supported
398 on Linux devices with Nvidia GPUs. MMORF uses a mixed computation
399 model, with only the most time-consuming components of the registration
400 algorithm being ported to the GPU. These include

401 • Image interpolation
402 • Cost, gradient and Hessian calculations
403 • Solving the system of linear equations for update steps

404 Of these, the Hessian calculation is by far the most computationally complex, with those of the regularisation and the rotational component of the
405 DTI cost functions being particularly burdensome. For these calculations,
406 GPU acceleration is on the order of 20-40x (depending on warp resolution,
407 image dimensions and image subsampling (Lange et al., 2020, Supplementary
408 Material: GPU Considerations and Code Profiling)).

410 MMORF has been designed from the outset for integration into the FSL
411 suite of neuroimaging tools (Jenkinson et al., 2012). As such, it expects
412 inputs in FSL convention. Specifically, affine matrices between input images
413 and the warp space should be in FLIRT format and DTI images should be in
414 dtifit format. MMORF output warp fields follow existing FSL conventions
415 and can therefore serve as drop-in replacements for FSL commands such as
416 applywarp.

417 3. Theoretical Differences Between Methods

418 With an understanding of the design considerations that went into MMORF,
419 we will now briefly describe how some of these choices compare to three ex-
420 isting registration tools, namely: FNIRT (Andersson et al., 2007), ANTs
421 Avants et al. (2008) and DR-TAMAS (Irfanoglu et al., 2016). These three
422 tools will then be used to validate and evaluate the performance of MMORF
423 in Section 4.

424 FNIRT was chosen as it is the predecessor to MMORF, and the cur-
425 rent nonlinear registration tool in FSL. The largest differences between these
426 methods are MMORF’s multimodal capabilities, regularisation, inverse-consistency
427 and GPU parallelisation. In terms of similarities, they share the same trans-
428 formation and bias field models, and very similar optimisation strategies at
429 low resolutions. At higher resolutions, FNIRT switches to a Scaled Conju-
430 gate Gradient algorithm (Møller, 1993), in contrast to MMORF’s MM ap-
431 proach. Finally, FNIRT performs a simultaneous optimisation of both warp
432 and bias fields, whereas these are optimised in a greedy, interleaved manner
433 by MMORF. This is because simultaneous optimisation results in a Hessian
434 without the regular diagonal structure, on which MMORF relies for efficient
435 GPU parallelisation. In practice, despite some similarities in design choices,
436 they perform very differently—even when MMORF is run, like FNIRT, in a
437 unimodal configuration. The resulting warps have a very different character,
438 which we largely attribute to the superior regularisation metric employed in
439 MMORF. MMORF’s inputs and output files are fully compatible with those

440 of FNIRT, and it can therefore serve as a drop-in replacement in FSL analysis
441 pipelines.

442 ANTs was chosen due to its consistently high performance, including in
443 dedicated registration comparisons (Klein et al., 2009). It has become a *de*
444 *facto* standard for nonlinear registration in much of medical imaging, and
445 serves as a benchmark against which to compare MMORF’s performance.
446 ANTs is a purely scalar registration method, although it can be applied to
447 multiple scalar input modalities simultaneously. ANTs uses a symmetric,
448 greedy approximation of large deformation diffeomorphic metric mapping
449 (LDDMM) known as symmetric normalisation (SyN). At each iteration, an
450 update step is composed with the current warp until convergence is reached..
451 If each update step is diffeomorphic, then the composition of all updates is
452 also diffeomorphic (apart from arithmetic floating point errors). Symmetry
453 is achieved by registering both reference and moving images to a mid-space
454 at each iteration. There are a number of similarity metrics implemented,
455 but we will limit ourselves to locally normalised cross-correlation (LNCC),
456 as this has proven to perform best in previous studies. ANTs regularisation
457 consists of simple Gaussian smoothing, which may be applied independently
458 to both update steps and to the final deformation field.

459 DR-TAMAS was chosen as it is currently the only other tool, to our
460 knowledge, which can simultaneously register both tensor and scalar modal-
461 ities in a single framework. It has also proven to match or exceed the best
462 performing diffusion registration tools currently available. As in MMORF,
463 finite-strain reorientation of the tensors is taken into consideration during
464 each update step (*i.e.*, it contributes to the gradient of the cost function,
465 and is not simply applied after each update step). In contrast to MMORF,
466 the DTI dissimilarity is divided into two separate cost functions—the trace
467 similarity (a rotationally invariant scalar), and the deviatoric tensor similar-
468 ity (sensitive to relative tensor orientation). For scalar inputs, DR-TAMAS
469 uses an LNCC cost function. The transformation model and optimisation
470 strategy are the same as ANTs, and therefore DR-TAMAS can be considered
471 as a truly multimodal variant of ANTs. In many ways, then, DR-TAMAS is
472 the natural alternative to MMORF.

473 There are clearly many other nonlinear registration tools against which we
474 could have compared, but we believe that these choices allow us to effectively
475 benchmark the relative performance of MMORF against those tools that are
476 most likely to be considered as an alternative by users.

477 4. Validation

478 Benchmarking a registration tool can be difficult to do well. Measures of
479 accuracy are often biased towards metrics based on the modalities that drove
480 the registration (Irfanoglu et al., 2016), and evaluations hence risk some de-
481 gree of circularity. Methods often perform best when evaluated using metrics
482 based on the same data that drove the registration (Irfanoglu et al., 2016),
483 which risks introducing a degree of circularity in these types of evaluations.
484 We have, therefore, endeavoured to perform a holistic evaluation of registra-
485 tion performance by including a range of structural-, diffusion-, functional-
486 and morphometry-derived metrics. Structural and diffusion metrics are the
487 most circular, since these are the modalities that drive the registration (ei-
488 ther individually or jointly) in all methods—however, they may highlight
489 both the value and the pitfalls of using the data you wish to analyse to
490 drive alignment. Functional metrics are based on a fully held-out modality
491 (that is not seen by any of the registration tools) and therefore serves as
492 the best proxy of true consistency of anatomical alignment across individ-
493 uals—since overfitting to a driving-modality (*e.g.*, unrealistically deforming
494 the brain to make structural images appear very similar) is likely to have a
495 negative effect on functional alignment (Coalson et al., 2018; Robinson et al.,
496 2018). Morphometry metrics (*e.g.*, measures of distortion that represent how
497 aggressively the images are deformed) are essential to contextualise and in-
498 terpret the accuracy/similarity metrics—that is, it is important to know how
499 aggressively a tool needs to deform an image to achieve a certain degree of
500 registration accuracy, with less distortion being preferred.

501 We chose the Human Connectome Project (HCP) young adult 1200 re-
502 lease (100 unrelated subjects subset) dataset as the basis for our testing
503 (Van Essen et al., 2012, 2013; Glasser et al., 2013). The HCP dataset con-
504 tains high quality T1w (0.7 mm isotropic), diffusion (1.25 mm isotropic) and
505 task-fMRI (2.0 mm isotropic) scans. This allows both unimodal T1w and
506 multimodal T1w + DTI registration to be conducted with the same dataset,
507 as well as evaluating registration metrics based on all three modalities. The
508 minimally preprocessed HCP data were used as far as possible (Glasser et al.,
509 2013), which includes motion and distortion correction with FSL topup and
510 eddy (Andersson et al., 2003; Andersson and Sotiroopoulos, 2016), and coreg-
511 istration of the diffusion data to T1w space. The task-fMRI data were,
512 however, reprocessed in subject T1w space with no smoothing (rather than
513 in MNI-152 space with 2.0 mm isotropic smoothing). The diffusion tensor

514 model was fit to the $b=1000\text{ s/mm}^2$ shell only using FSL dtifit.

515 The Oxford-MultiModal-1 (OMM-1)² template was used as the reference
516 space for all registrations (Arthofer et al., 2021, 2022). This template was
517 constructed from 240 UK Biobank (Miller et al., 2016) subjects, and has
518 both T1w and DTI volumes that are intrinsically co-registered. The OMM-1
519 therefore provides a common space in which to compare methods that either
520 use T1w images only, or a combination of T1w and DTI images to drive the
521 registration.

522 We calculate two structural accuracy metrics based on pairwise label sim-
523 ilarity. The first is a measure of overlap (specifically, the Jaccard coefficient)
524 and the second a measure of distance (specifically, the Hausdorff distance)
525 of automatically generated cortical and subcortical labels for each pair of
526 subjects following resampling to template space (see Section 4.2 for details).
527 The Jaccard coefficient for two regions A and B is defined as:

$$\frac{A \cap B}{A \cup B}, \quad (14)$$

528 and ranges from 0 for no overlap to 1 for perfect overlap. The Hausdorff
529 distance is the maximum distance between the surfaces of the regions being
530 compared. The labels are derived from the T1w images, and therefore we
531 expect them to favour scalar registration methods that can match the T1w
532 images without being penalised if that reduces DTI similarity. We have
533 avoided any simple intensity-based or tissue-type overlap metrics, as these
534 are known to correlate poorly with true anatomical consistency (Rohlfing,
535 2012).

536 We then calculate three diffusion accuracy metrics that compare the sim-
537 ilarity of the template DTI modality with each subject's DTI data after
538 resampling to template space. Each subject is compared voxelwise to the
539 template, and the average across all voxels within a mask is taken as the
540 overall metric for that subject. Overall tensor similarity (OVL, Equation 15)
541 is the first metric, and balances directional and magnitude similarity between
542 tensors and is a good general indicator of the similarity between two tensors.
543 Linear-shape weighted V1 similarity (CLV1, Equation 17) is the second met-
544 ric, and is defined as the inner product of the first eigenvector of the diffusion
545 tensor (V1) from the template with V1 from the warped subject, weighted

²DOI 10.17605/OSF.IO/S9GE4

546 by the coefficient of linear shape (CL) for the template tensor (*i.e.*, how sim-
547 ilar V1 is, weighted by how informative V1 is). Planar-shape weighted V3
548 (CPV3, Equation 19) is the third metric, and is defined as the inner product
549 of the third eigenvector of the diffusion tensor (V3) from the template with
550 V3 from the warped subject weighted by the coefficient of planar shape (CP)
551 for the template tensor (*i.e.*, how similar V3 is, weighted by how informative
552 V3 is). See Figure 1 for a visual depiction of CL and CP maps of the OMM-
553 1 template. In areas where CL is large, the tensor shape is largely prolate
554 (cigar-shaped), and therefore the direction of maximum diffusivity (*i.e.*, V1)
555 is well defined. In areas where CP is large, the tensor shape is largely oblate
556 (plate-shaped), and therefore the direction of minimum diffusivity (*i.e.*, V3)
557 is well defined. CLV1 and CPV3 specifically probe how well orientational in-
558 formation has been aligned, which is most relevant in white matter regions.
559 We expect those tools that consider rotational information (*i.e.*, DTI data)
560 during registration to perform best in these metrics.

561 The fMRI accuracy metric used is task-fMRI cluster mass. This measures
562 how consistently the registration tools are able to align those regions of the
563 brain that are significantly activated (or deactivated) when performing a task.
564 This assumes a general correspondence between brain structure and function,
565 but this need not be exact. The benefit of this metric is that it is entirely
566 independent of the modalities driving registration (*i.e.*, T1w and DTI) and,
567 therefore, there is little to no circularity in its interpretation, which cannot
568 be said for the previous metrics.

569 Finally we calculate both size ($|\mathbf{J}|$) and shape (CVAR—see Section 4.2.5
570 for definition) distortion metrics to evaluate how much each tool has had
571 to deform the subject’s images to match the template. For any given level
572 of accuracy (*i.e.*, the preceding metrics) a smaller amount of distortion is
573 usually preferred, as this indicates that the registration method is changing
574 the original data as little as possible.

575 4.1. Ethics

576 All human imaging data used in this work are part of the open access
577 Human Connectome Project Young Adult (HCP-YA) dataset. The data are
578 pseudonymised and identifiable visual features, such as the face and ears, are
579 obscured. Written informed consent to share the data “broadly and openly”
580 was obtained for all participants by the original HCP-YA researchers (Elam
581 et al., 2021).

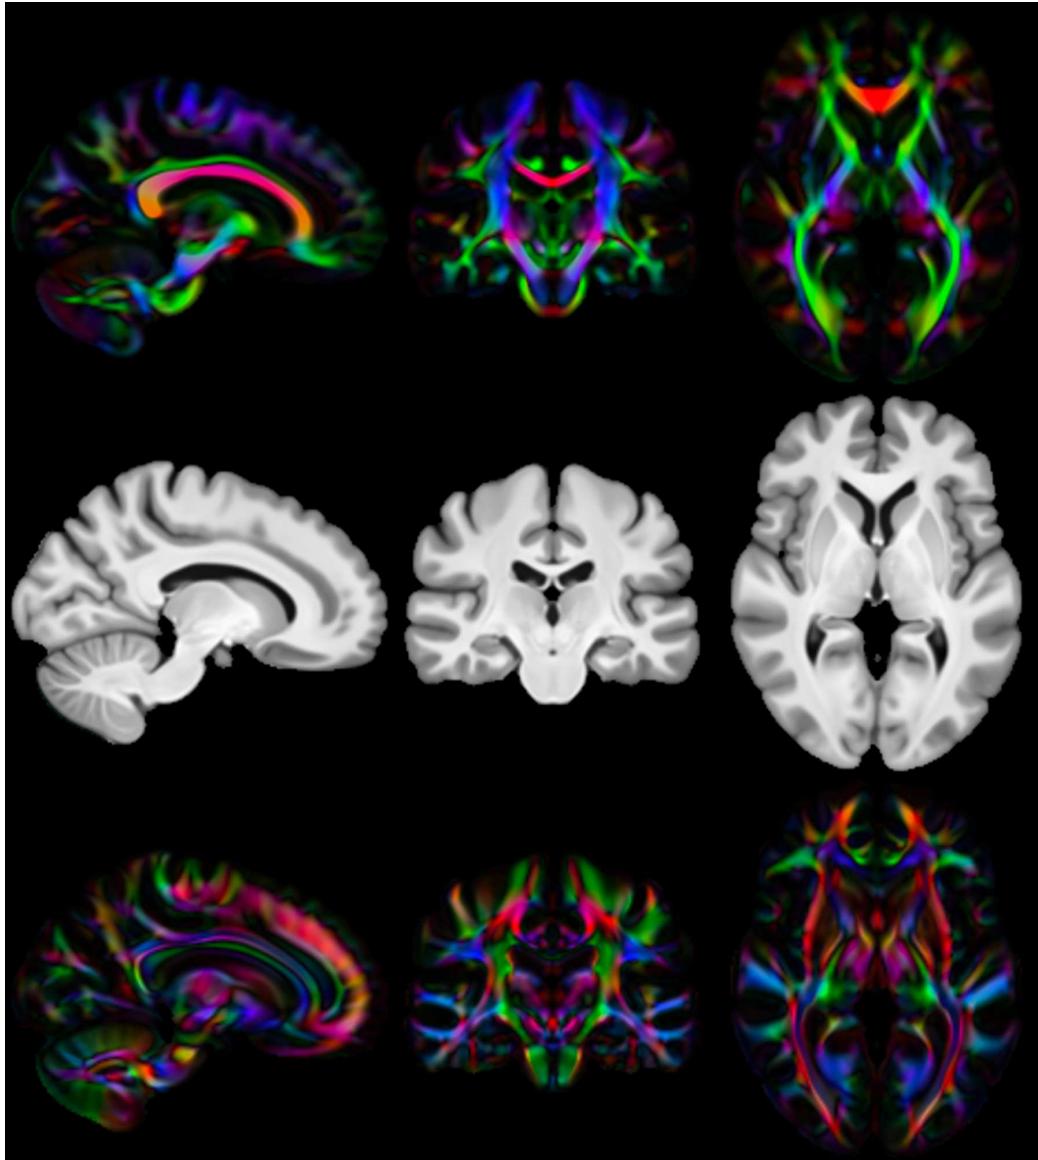


Figure 1: **Visualisation of linear and planar shape coefficients in OMM-1 space:** From top to bottom: Linear coefficient (CL) map of the template, T1w template for reference, planar coefficient (CP) map of the template. Images are directionally encoded colour maps of V1 and V3 respectively. Green = Anterior-Posterior, Red = Left-Right, Blue = Inferior-Superior. Both CL and CP values are highest in the white matter, but cover complementary regions therein. CLV1 is most sensitive to how well the primary diffusion direction is matched in voxels where diffusion occurs parallel to that axis only. CPV3 is most sensitive to how well the tertiary diffusion direction is matched in voxels where diffusion occurs within a plane defined by the primary and secondary diffusion directions. Together, these represent the voxels where tensor orientation can be reliably described by a single direction.

582 *4.2. Methods*

583 Each of the following steps was performed independently for each regis-
584 tration method.

585 *4.2.1. Registration*

586 All registration was performed to the OMM-1 template. This template
587 contains both T1w and DTI volumes with isotropic resolution of 1 mm. Since
588 both modalities were jointly aligned during template creation, they are intrin-
589 sically spatially consistent at the voxel level. It is therefore a good choice for
590 registration with both unimodal and multimodal tools. Each tool was used
591 to register all of the 100 unrelated HCP subjects to the template. FLIRT,
592 FNIRT and ANTs used the T1w image only, while MMORF and DR-TAMAS
593 used both the T1w and DTI images. The FLIRT affine matrix was used to
594 initialise both FNIRT and MMORF, whereas ANTs and DR-TAMAS used
595 their own affine registration methods. FNIRT and MMORF were run with
596 custom configurations identified to empirically produce good results. ANTs
597 and DR-TAMAS were run with slightly modified configurations that were
598 found to improve registration accuracy over the defaults. All methods were
599 run using a multi-resolution pyramid approach to a final warp resolution of
600 1 mm isotropic.

601 *4.2.2. T1w FreeSurfer Label Overlap*

602 Automatically segmented subcortical (ASEG atlas, (Fischl et al., 2002,
603 2004)) and cortical (Destrieux 2009 atlas, (Destrieux et al., 2010)) parcella-
604 tions for each subject were warped into template space. Jaccard coefficients
605 (measuring label overlap) and Hausdorff distances (measuring the maximum
606 error in label boundaries) of the corresponding warped parcellations were
607 calculated for every possible pair of subjects. A pairwise approach was used
608 because there are no target labels in template space. The average coefficient
609 for each parcellation was then calculated across pairings for each subject,
610 resulting in 100 values for each parcellation.

611 *4.2.3. DTI Similarity*

612 Combined affine and nonlinear (apart from FLIRT) warp fields were used
613 to resample each subject's DTI volume into template space. Resampling of
614 tensors was performed using the FSL tool `vecreg`, which includes preserva-
615 tion of principle direction (Alexander et al., 2001) reorientation, which is the
616 most accurate way of accounting for the rotational effect of the warps when

617 resampling DTI data. Tensors were then decomposed into three Eigenvalue-
618 Eigenvector pairs (L1/2/3 and V1/2/3 respectively). OVL, CLV1 and CPV3
619 were then calculated between each subject's warped DTI volume and the
620 template. A pairwise approach was not necessary here since the DTI tem-
621 plate volume itself acts as the target. CL and CP weighting coefficients
622 (Alexander et al., 2000) were generated from the template DTI volume only.
623 V1 and V3 similarity was calculated as the magnitude of the dot product
624 of template and warped subject vectors in each template voxel. Similarity
625 metrics were calculated per voxel according to the following formulae:

$$OVL = \frac{\sum_{i=1}^3 L_{i_r} L_{i_m} (V_{i_r}^\top V_{i_m})^2}{\sum_{i=1}^3 L_{i_r} L_{i_m}} \quad (15)$$

$$CL = \frac{L_{1_r} - L_{2_r}}{L_{1_r} + L_{2_r} + L_{3_r}} \quad (16)$$

$$CLV1 = CL |V_{1_r}^\top V_{1_m}| \quad (17)$$

$$CP = \frac{2 (L_{2_r} - L_{3_r})}{L_{1_r} + L_{2_r} + L_{3_r}} \quad (18)$$

$$CPV3 = CP |V_{3_r}^\top V_{3_m}| \quad (19)$$

626 where the subscripts r and m represent the reference and resampled moving
627 image respectively. All metrics were calculated within the template brain
628 mask only.

629 4.2.4. *tfMRI Cluster Mass*

630 The HCP task battery (Barch et al., 2013) consists of 7 tasks from which
631 86 contrasts are derived, and used to generate contrast of parameter esti-
632 mate (COPE) images in subject T1w space. For each subject, the 86 COPE
633 images were resampled into template space. FSL Randomise (Winkler et al.,
634 2014) was then used to perform a group-level, non-parametric, ordinary-least-
635 squares, random-effects, one-group *t*-test on the mean COPE image (across
636 subjects) for each contrast. The results of the group-level processing are 86
637 *t*-statistic maps/images and 86 family-wise-error (FWE) corrected *p*-value
638 maps/images (one per contrast). The *t*-statistic map represents the group
639 activation for a particular contrast, where more accurately aligned activa-
640 tions lead to higher *t*-statistics. The *p*-value map represents the statistical
641 significance of the *t*-statistic at each voxel.

642 Cluster mass was then calculated as follows. Thresholding was applied
643 to the FWE-corrected p -value map (at $p < 0.05$) for each contrast, to give
644 a binary mask of significantly activated regions. This mask was then ap-
645 plied voxel-wise to the t -statistic map. The masked t -statistic map was then
646 multiplied voxel-wise with a grey matter mask of the template, and summed
647 across all voxels. The summed value is what we refer to as cluster mass, and
648 is a single number per contrast and registration method. Cluster mass is,
649 therefore, increased by both higher t -statistics and better alignment of grey
650 matter. It can be calculated independently for each registration method,
651 simplifying between-method comparisons

652 *4.2.5. Distortion*

653 Two measures of distortion are considered, both evaluated within the
654 template brain mask.

655 The first is the 5th to 95th percentile range of the log-Jacobian determinant
656 ($\log |\mathbf{J}|$). This is a measure of volumetric distortion. Since the histogram
657 of $\log |\mathbf{J}|$ tends to be centred around zero, the mean is uninformative and,
658 therefore, a range is more useful. This a fairly robust measure of the extent
659 to which the local effect of the warp field is to expand/contract voxels (and
660 is equally sensitive to both).

661 The second is the average cube-volume aspect ratio (CVAR, (Smith and
662 Wormald, 1998)), defined as:

$$663 CVAR = \sqrt[3]{\frac{s_{i-\max}^3}{s_1 \times s_2 \times s_3}} \quad (20)$$

664 where s_i are the singular values of the local Jacobian matrix. CVAR repre-
665 sents the cube-root of the ratio of the volume of the smallest regular cube
666 which can fully enclose the cuboid, to the cuboid's own volume. Alterna-
667 tively, for a deformation one can equally well define it as the cube-root of the
668 ratio of the largest Jacobian singular value cubed, to the Jacobian deter-
669 minant. This is a pure shape-distortion measure that is invariant to volumetric
670 changes. For a perfect cube its value is 1, and it is greater than 1 for any
671 shape where one or more sides of the cube are a different length to the others.
672 Since it is always ≥ 1 , the mean across voxels is a meaningful measure of the
673 extent to which the local effect of the warp field is to alter the original shape
674 of the underlying voxels.

675 Between them, these two measures present a good summary of the extent
676 to which the warp is distorting the anatomy of the brain to achieve a given

Table 1: **Summary of results for all domains.** Structural (FreeSurfer labels): median Jaccard index (JI) and Hausdorff distance (HD) across subjects for both subcortical and cortical labels. Diffusion (DTI): median overall tensor similarity (OVL), linear coefficient-weighted V1 similarity (CLV1) and planar coefficient-weighted V3 (CPV3) similarity. Functional (tfMRI): Total cluster mass (CM) for all contrasts. Distortion: median 5th to 95th log-Jacobian determinant range ($|\mathbf{J}|$) and cube-volume aspect ratio (CVAR) across subjects. The best performing method in each metric is highlighted.

	FreeSurfer Labels				DTI			tfMRI	Distortion	
	Subcort		Cort		OVL	CLV1	CPV3	CM	$ \mathbf{J} $	CVAR
	JI	HD	JI	HD						
FLIRT	0.449	5.794	0.197	10.797	0.669	0.802	0.755	2.967e+7	-	-
FNIRT	0.606	5.160	0.409	9.329	0.776	0.873	0.854	3.718e+7	1.858	1.496
ANTs	0.652	4.707	0.397	9.147	0.802	0.886	0.867	3.736e+7	1.197	1.366
DR-TAMAS	0.640	5.287	0.373	9.406	0.817	0.896	0.869	3.906e+7	1.390	1.374
MMORF	0.664	4.891	0.399	9.102	0.825	0.900	0.872	3.917e+7	1.202	1.353

674 degree of registration accuracy. In general, lower distortion is for a given
 675 level of accuracy is preferred. When both distortion and accuracy increase
 676 then a more nuanced interpretation is required based on the relative changes
 677 in each.

678 4.3. Results

679 The results across all domains and metrics are summarised in Table 1.
 680 More detailed descriptions with accompanying figures are presented in the
 681 subsections that follow. Many figures utilise raincloud plots (Allen et al.,
 682 2021) that simultaneously show the raw data, a box and whiskers plot, and a
 683 density estimate. In all cases the interpretation of the box and whisker part
 684 of the plots is the same—the box shows the quartiles (25th, 50th and 75th
 685 percentiles), the whiskers extend to the final non-outlier datapoint, and the
 686 diamonds show outliers (any point more than 1.5 \times the interquartile range
 687 from the edges of the box).

688 4.3.1. T1w FreeSurfer Labels

689 The Jaccard index and Hausdorff distance results are presented in Figures
 690 2 and 3 respectively. In all cases, nonlinear registration leads to a major
 691 improvement in both metrics. MMORF and ANTs produce the best Jaccard
 692 index results in the subcortex, with FNIRT narrowly outperforming them in
 693 cortical regions. Hausdorff distance performance is slightly better in both
 694 subcortical and cortical regions for MMORF and ANTs. It is worth noting
 695 that subcortical contrast is relatively poor in T1w images from the HCP

696 dataset, and certain structures (the left pallidum in particular) are poorly
697 segmented in a number of subjects. This is the cause of the relatively heavy
698 tails towards low Jaccard coefficients in Figure 2a.

699 *4.3.2. DTI Similarity*

700 OVL, CLV1 and CPV3 similarity results are presented in Figures 4 to 6 re-
701 spectively. In all cases, the trend is the same: affine registration with FLIRT
702 performs worst followed by FNIRT, ANTs, DR-TAMAS, and MMORF (in
703 that order).

704 *4.3.3. tfMRI Cluster Mass*

705 Cluster mass results are presented in the form of percentage difference
706 plots, and warrant some guidance on their interpretation. Each point on the
707 plot represents one contrast. When comparing method *A* to method *B*, the
708 x-axis represents the cluster mass of method *A*. The *x*-axis is log-transformed
709 to account for the large range in cluster mass across contrasts. The *y*-axis
710 represents the percentage improvement in cluster mass by method *A* over
711 method *B*. Therefore, a point at location $(x, y) = (10, 10)$ represents a con-
712 trast with a cluster mass of $\exp(10)$ for method *A*, and a 10 % improvement
713 in cluster mass when using method *A* over method *B*. Similarly, a point at
714 location $(x, y) = (15, -5)$ represents a contrast with a cluster mass of $\exp(15)$
715 for method *A*, and a 5 % reduction in cluster mass when using method *A* over
716 method *B*. By choosing percentage difference for the *y*-axis we remove the
717 bias towards large cluster mass contrasts, which would otherwise dominate
718 if *y* were instead the simple difference in cluster mass between methods.

719 Cluster mass comparisons for MMORF vs FLIRT, FNIRT, ANTs and
720 DR-TAMAS are presented in Figures 7 to 10 respectively. MMORF produces
721 a large improvement across all contrasts compared to FLIRT, confirming
722 that the CM metric is sensitive to improved registration accuracy. MMORF
723 improves CM across most contrasts compared to FNIRT and, to a slightly
724 lesser extent, ANTs. MMORF and DR-TAMAS produce very similar results
725 for contrasts with high CM (towards the right of the x-axis), but MMORF
726 performs consistently better for small and medium CM contrasts.

727 *4.3.4. Distortion*

728 Comparisons of the 5th to 95th percentile Jacobian determinant range
729 (volumetric distortion) and the average CVAR (shape distortion), which are
730 both evaluated only within a brain mask, are presented in Figures 11 and

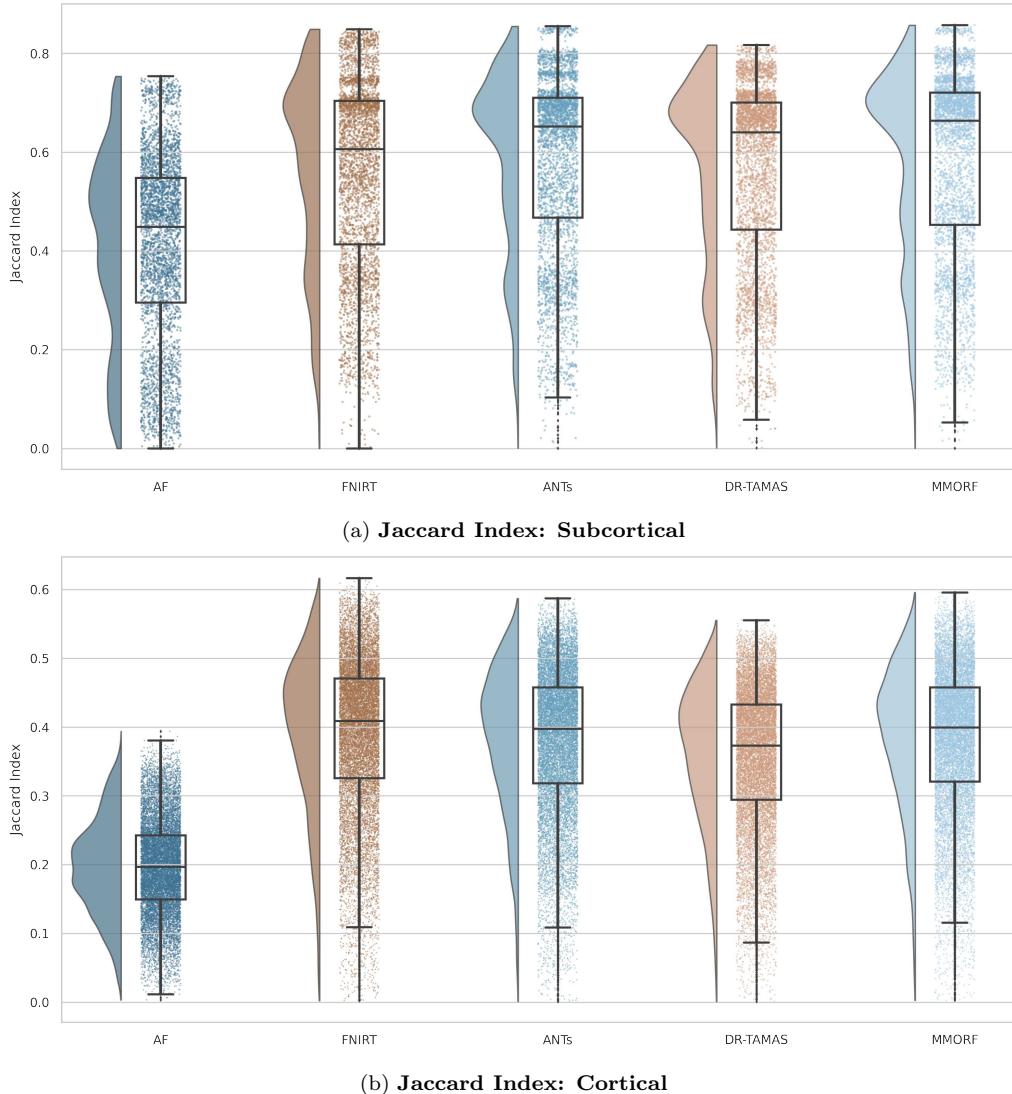
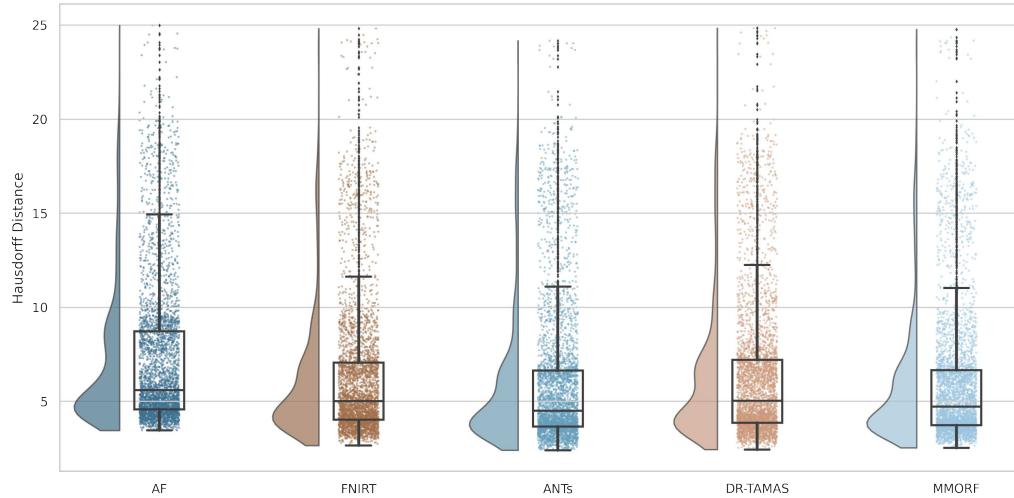
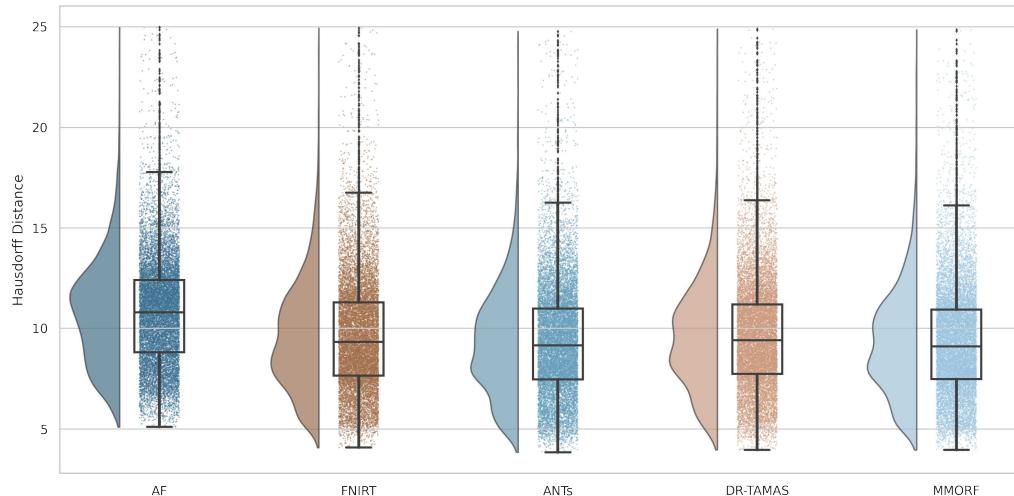


Figure 2: **Subcortical (a) and cortical (b) Jaccard indices for FreeSurfer segmentation overlaps:** All nonlinear methods improve over affine only, with the greatest improvement being in the cortex. Across all labels, MMORF and ANTs perform similarly (and best), with only FNIRT slightly outperforming them in the cortex.



(a) Hausdorff Distance: Subcortical



(b) Hausdorff Distance: Cortical

Figure 3: Subcortical (a) and cortical (b) Hausdorff distances for FreeSurfer segmentation overlaps: All nonlinear methods improve over affine only, with larger improvements evident in the cortex. Performance is very similar across all methods, with ANTs and MMORF slightly outperforming the others.

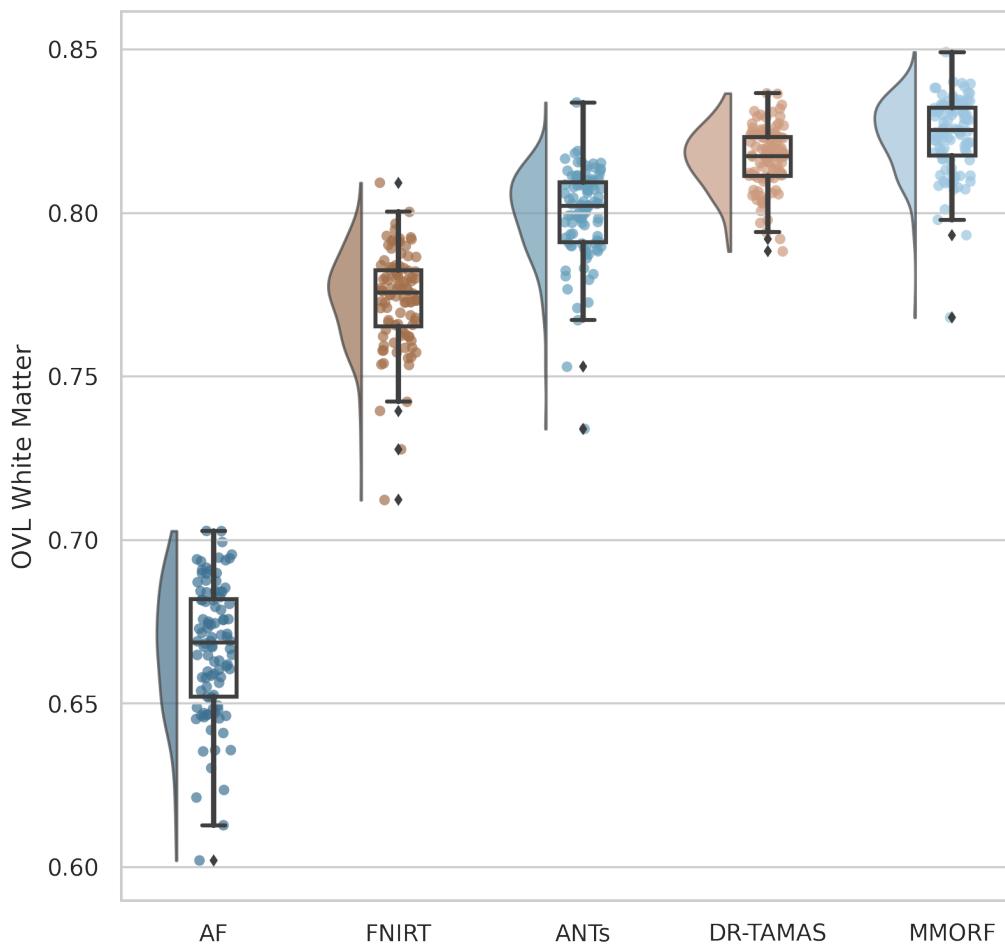


Figure 4: **Overall tensor similarity (OVL):** Calculated within a mask of the template white matter, defined as $FA > 0.2$. All nonlinear methods improve over affine only, with those methods that include DTI data in the registration (MMORF and DR-TAMAS) outperforming the T1w-only methods (FNIRT and ANTs). MMORF performs best overall.

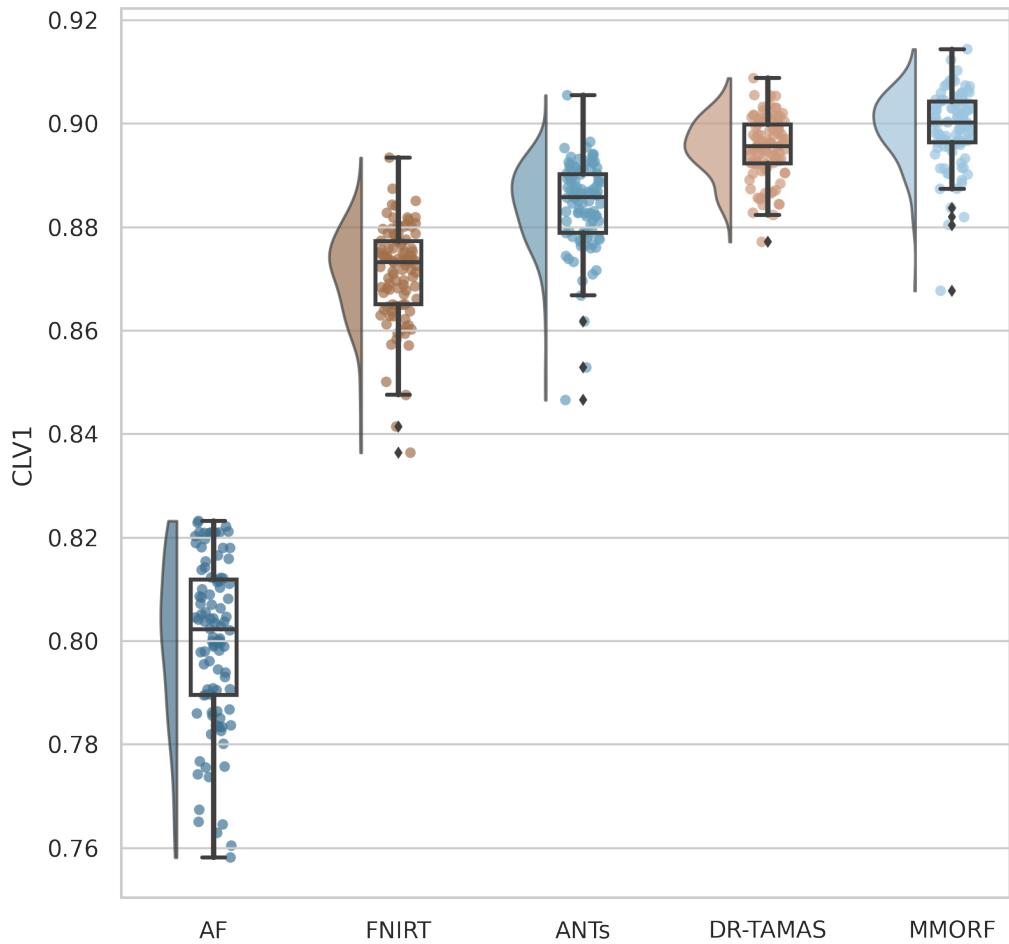


Figure 5: **Linear shape weighted V1 similarity (CLV1):** All nonlinear methods improve over affine only, with those methods that include DTI data in the registration (MMORF and DR-TAMAS) outperforming the T1w-only methods (FNIRT and ANTs). MMORF performs best overall.

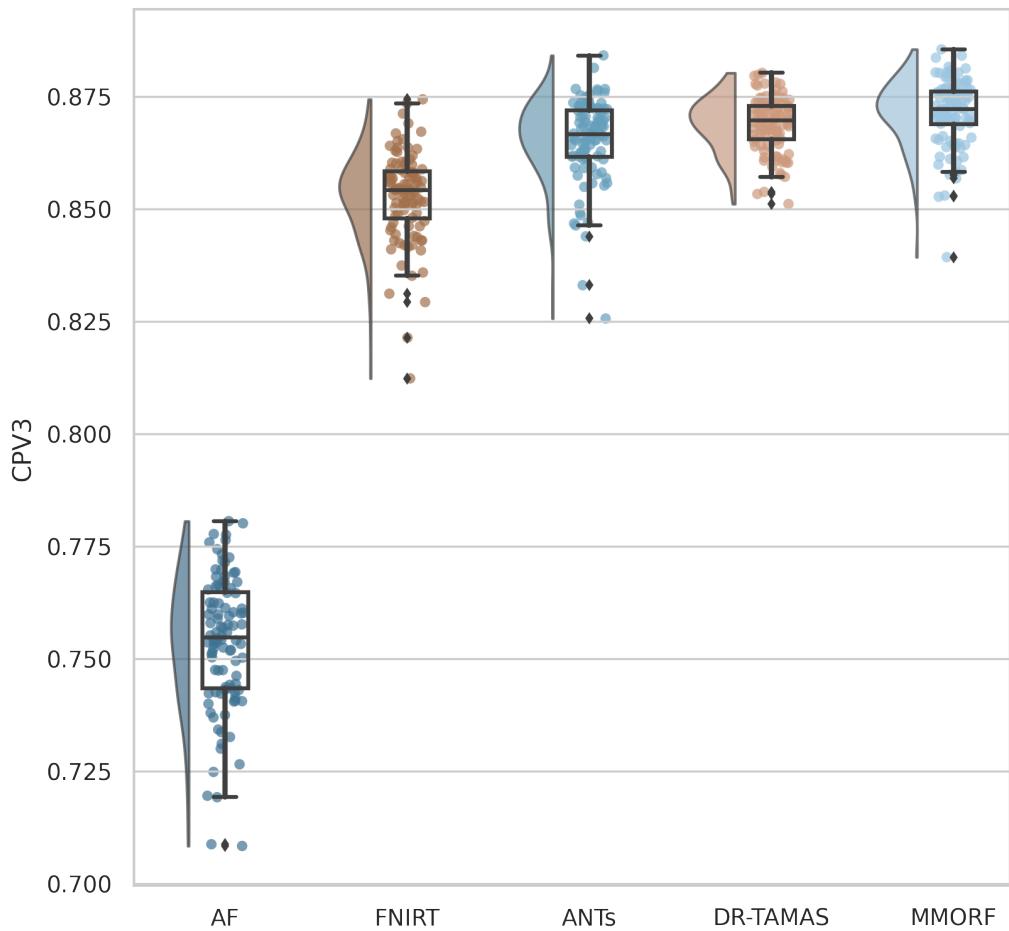


Figure 6: **Planar shape weighted V3 similarity (CPV3):** All nonlinear methods improve over affine only, with those methods that include DTI data in the registration (MMORF and DR-TAMAS) outperforming the T1w-only methods (FNIRT and ANTs). MMORF performs best overall.

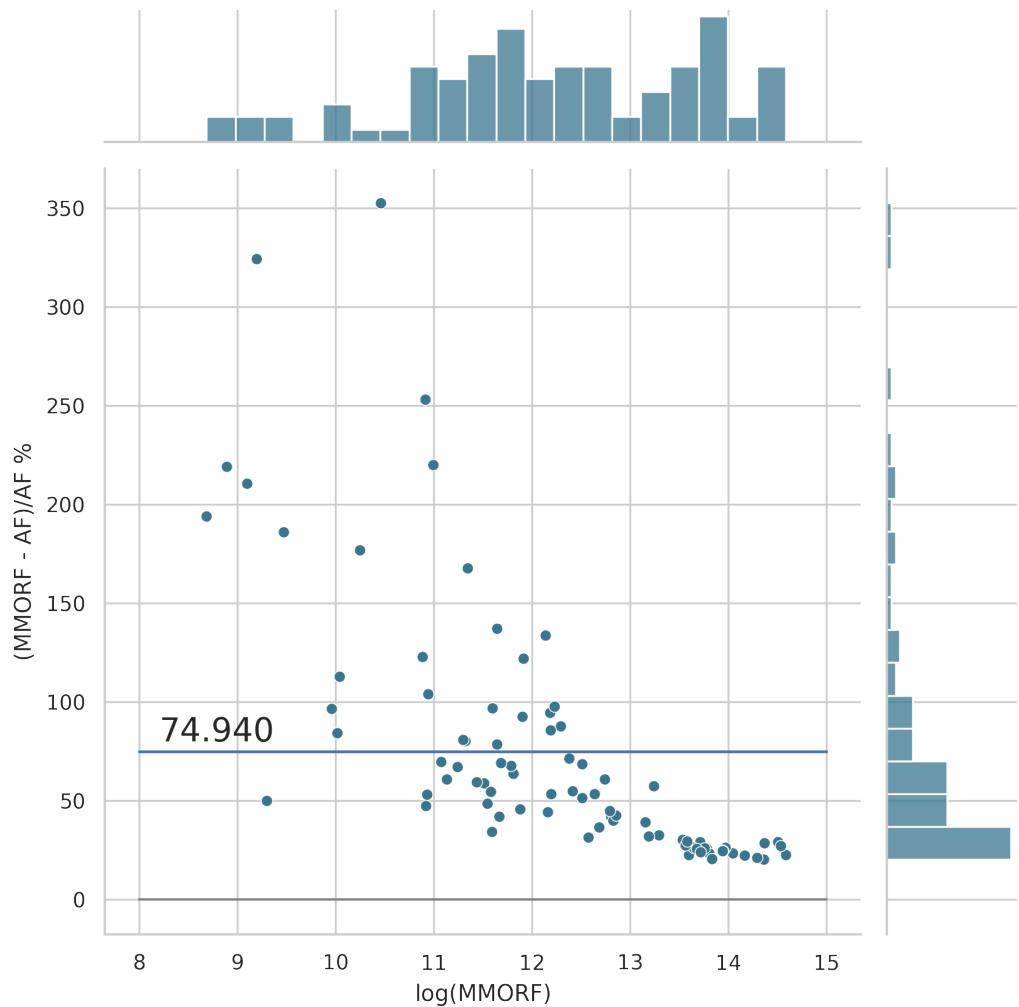


Figure 7: **MMORF vs FLIRT cluster mass:** Across clusters of all size, MMORF (nonlinear) outperforms FLIRT (linear), with an average improvement in cluster mass of $\approx 75\%$. This is not surprising, but serves to demonstrate that cluster mass is sensitive to registration accuracy.

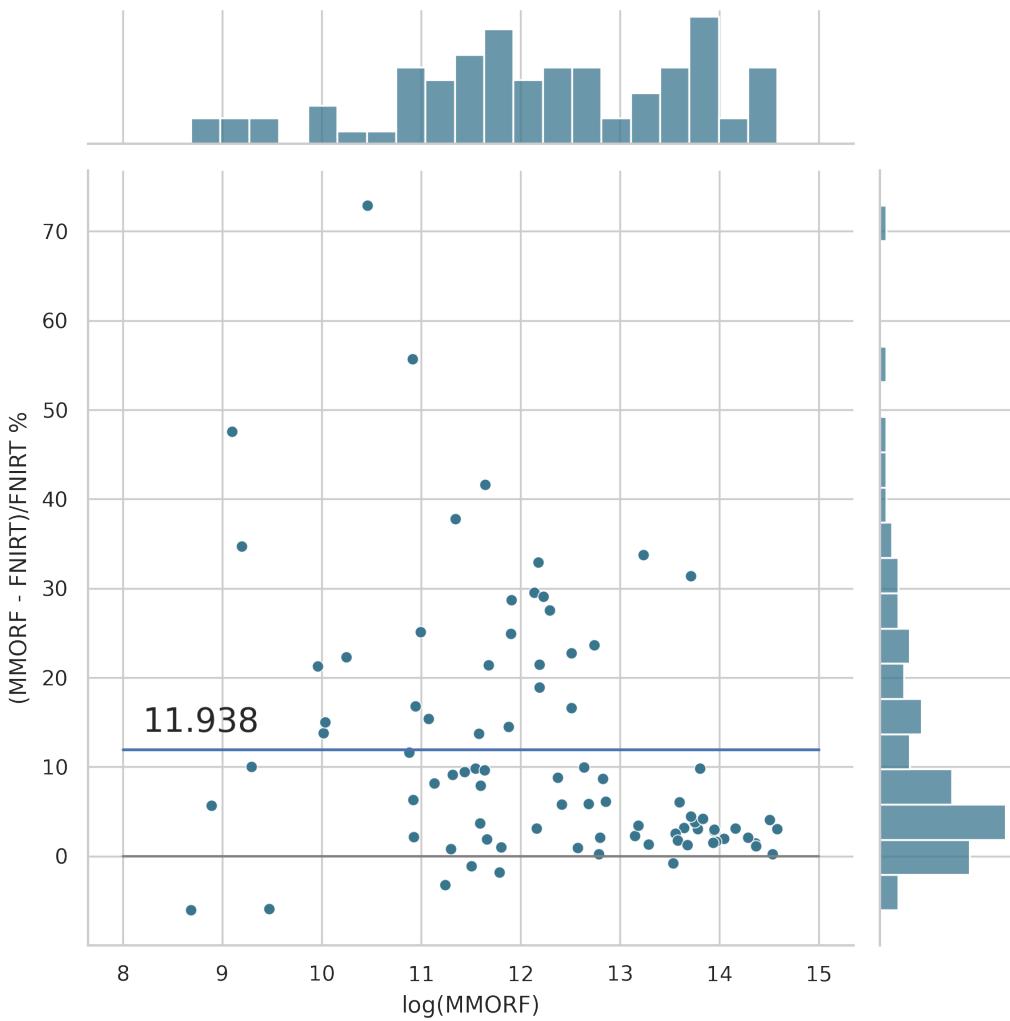


Figure 8: **MMORF vs FNIRT cluster mass:** MMORF outperforms FNIRT with an average improvement in cluster mass of $\approx 12\%$. The largest improvements are in the mid-sized clusters.

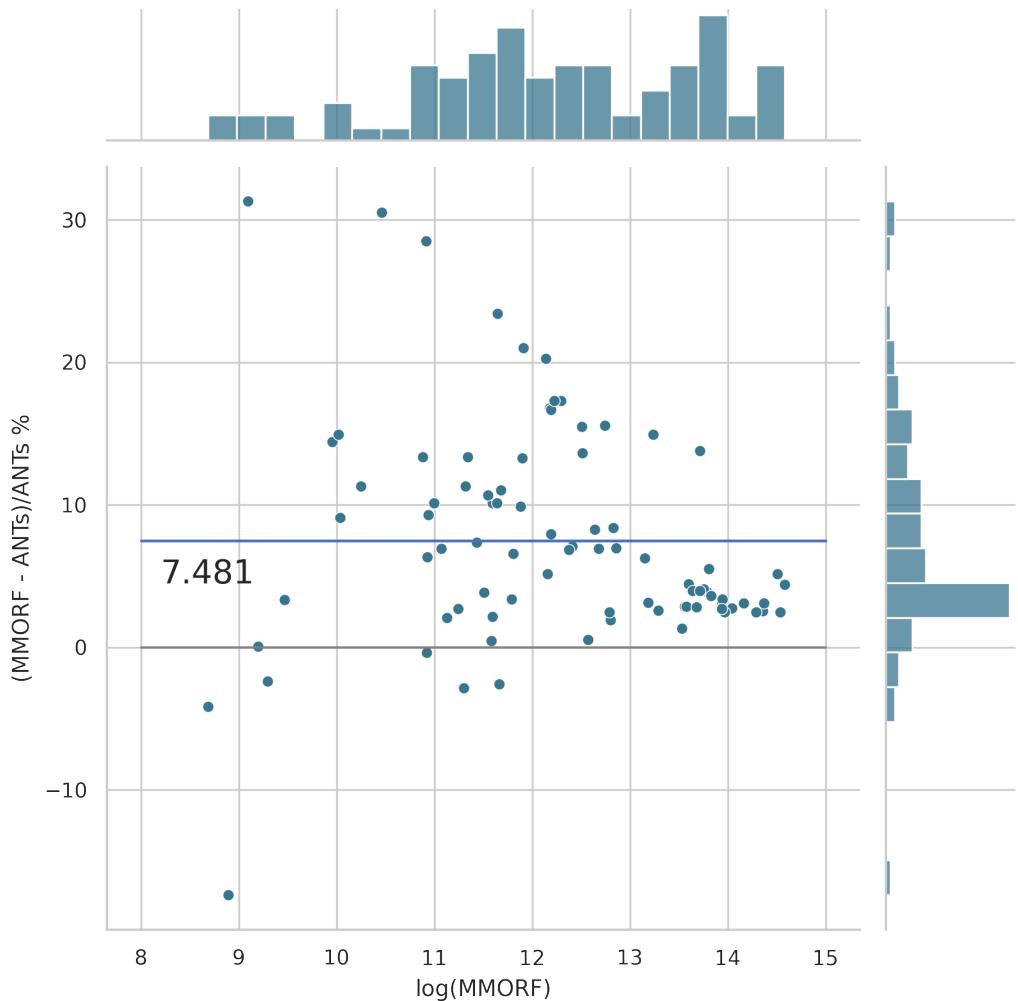


Figure 9: **MMORF vs ANTs cluster mass:** MMORF outperforms ANTs with an average improvement in cluster mass of $\approx 7.5\%$. The largest improvements are in the mid-sized clusters.

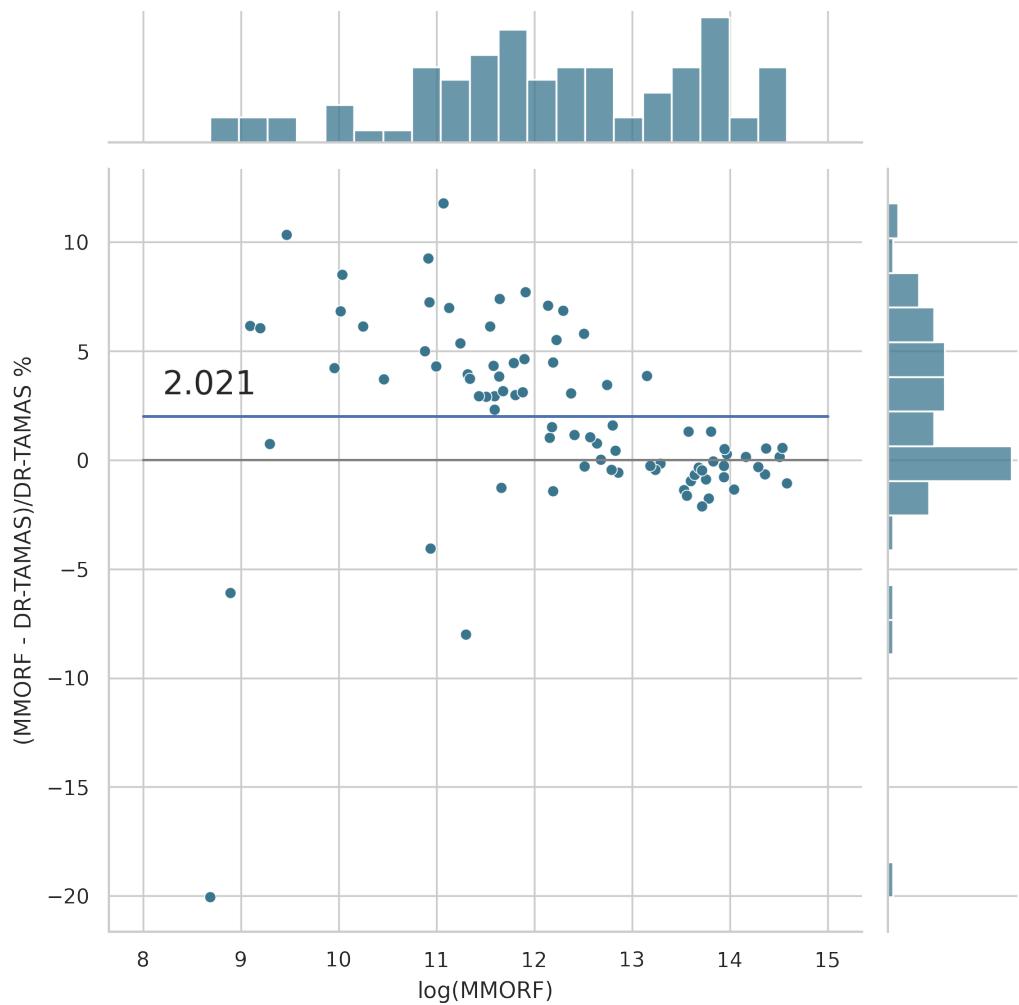


Figure 10: **MMORF vs DR-TAMAS cluster mass:** MMORF outperforms DR-TAMAS with an average improvement in cluster mass of $\approx 2\%$. The largest improvements are in the mid-sized clusters.

731 12 respectively. FNIRT shows significantly more distortion than the other
732 methods. ANTs and MMORF display similar levels of volumetric distortion,
733 that are lower than DR-TAMAS. MMORF displays the least shape distor-
734 tion, followed by ANTs and DR-TAMAS. The difference between ANTs and
735 DR-TAMAS is likely due to larger deformations within the white matter with
736 DR-TAMAS, due to diffusion information driving the registration harder in
737 those regions.

738 *4.4. Summary*

739 Of the methods tested, MMORF is the most consistently high-performing
740 across the full range of evaluation metrics.

741 It might be expected that the T1w-only-driven registration methods would
742 perform well when evaluated with a T1w-derived similarity metric (particu-
743 larly in the cortex) and, indeed, both FNIRT and ANTs perform very well in
744 the label overlap metric in this region. It is also possible that the inclusion
745 of DTI information during registration might negatively affect such a metric
746 and, again, we do indeed see that DR-TAMAS performs relatively poorly in
747 the cortex, despite good subcortical performance. MMORF, however, does
748 not seem to suffer in the same way—performing on par with ANTs both
749 cortically and subcortically. We may, therefore, conclude that MMORF is
750 a good choice of method, on par with ANTs, when a structurally-derived
751 segmentation comparison is the type of study for which registration is being
752 employed.

753 The value of including DTI information in the registration is clear, if
754 unsurprising. MMORF and DR-TAMAS noticeably outperform both FNIRT
755 and ANTs across all DTI similarity metrics. Of the multimodal registration
756 methods, MMORF has the advantage over DR-TAMAS across all metrics.
757 The OVL metric shows that, within the white matter, both the shape and
758 the size of the tensors are better matched by MMORF. The CLV1 and CPV3
759 metrics show that the most informative directions of the tensor are also best
760 aligned by MMORF.

761 We believe the tfMRI results to be the most unbiased assessment of
762 registration accuracy presented here, since it is evaluated using a modal-
763 ity that was never seen by any of the registration methods being tested.
764 Since MMORF outperforms all other methods under test, this is the clearest
765 indication of its high registration accuracy. It is notable that FNIRT, the
766 best performing method in terms of cortical label overlaps, does not perform
767 as well in this (also cortical) metric. Similarly, DR-TAMAS is the closest

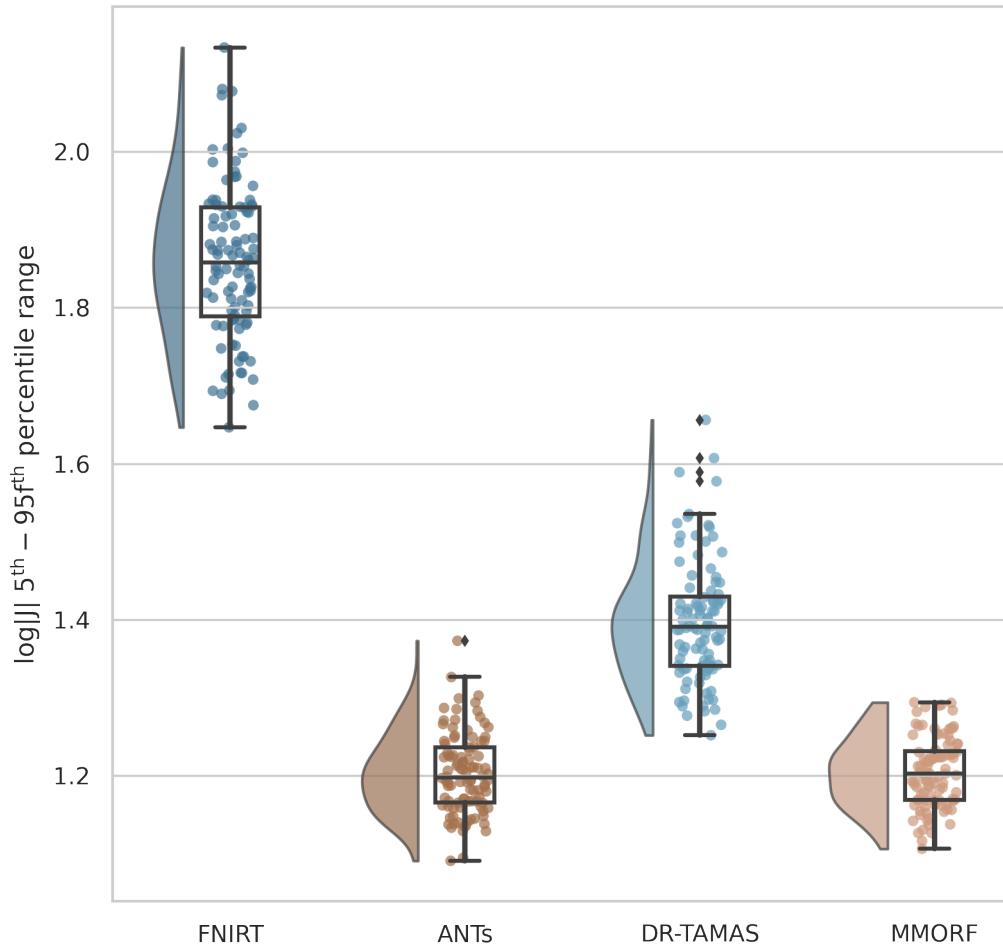


Figure 11: **5th to 95th percentile Jacobian determinant range:** MMORF and ANTs have the lowest, and very similar, levels of volumetric distortion. This is despite the fact that the MMORF warps are also trying to align the DTI information in the white matter, which would be expected to increase the amount of distortion (as can be seen has happened for DR-TAMAS). FNIRT shows the largest amount of distortion.

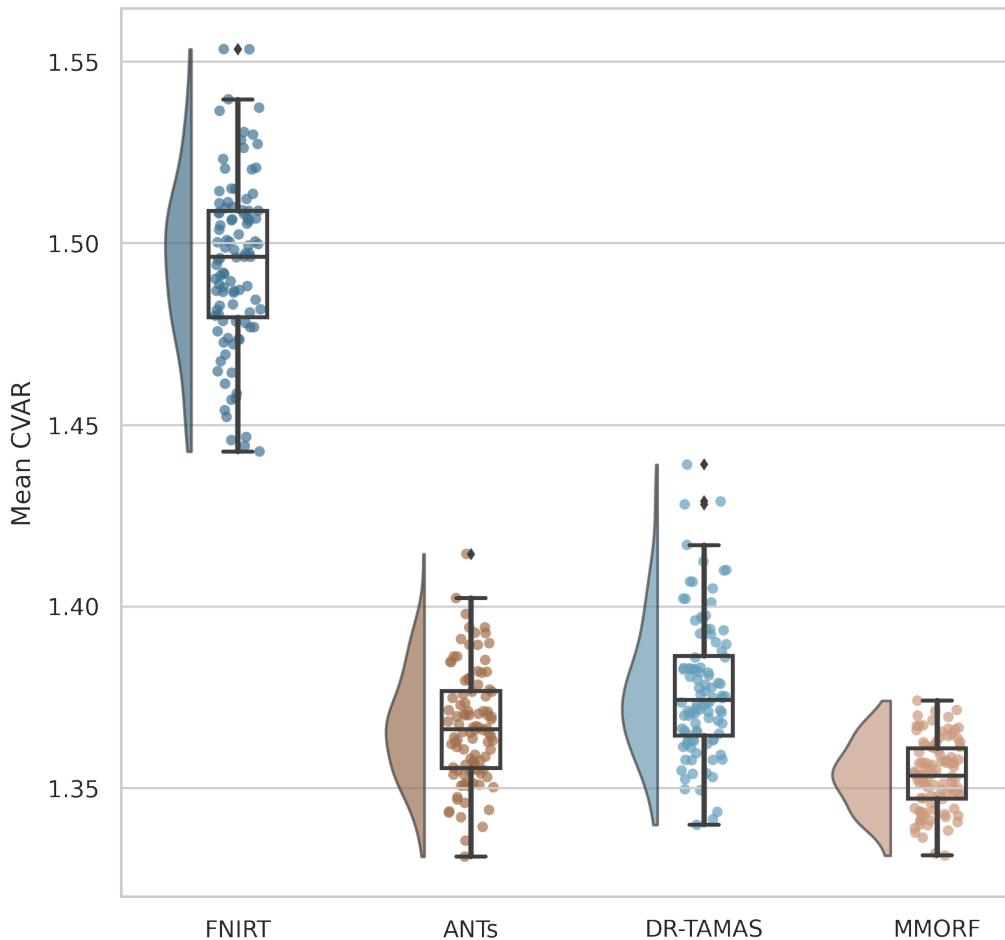


Figure 12: **Mean cube-volume aspect ratio (CVAR) shape distortion:** MMORF produces the lowest level of shape distortion on average, followed by ANTs, DR-TAMAS and FNIRT. As with the volumetric distortion, this is in spite of the fact that the MMORF warps are also trying to align the DTI information in the white matter, which would be expected to increase the amount of distortion (as can be seen has happened for DR-TAMAS). This also demonstrates that MMORF's low levels of volumetric distortion (see Figure 11) do not come at the cost of increased shape distortion, which can be observed when using a regularisation that penalises only the Jacobian determinant (Lange et al., 2020).

768 performing method to MMORF in the tfMRI evaluation, but is the poorest
769 performing method in cortical label overlap. This is a clear example of
770 the benefit in not only considering the more circular evaluation metrics (i.e.,
771 those derived from the same modalities that are driving the registration)
772 when comparing registration tools.

773 Finally, in terms of distortion, MMORF demonstrates levels that are simi-
774 lar to or lower than those of ANTs, despite matching it in label overlap and
775 bettering it in both the DTI and tfMRI metrics. DR-TAMAS produces more
776 distortion on average than its sibling method, ANTs, which is not unexpected
777 given that DR-TAMAS is also being driven by DTI information in the white
778 matter. FNIRT produces larger distortions than the other nonlinear meth-
779 ods. Without MMORF for comparison, this might easily be interpreted as
780 the inevitable consequence of driving a small deformation framework method
781 very hard to maximise accuracy metrics. However, given that MMORF and
782 FNIRT have the same transformation model, this instead highlights the value
783 of using a more biologically plausible regularisation model within the small-
784 deformation framework.

785 5. Discussion

786 The primary objective in developing MMORF was to exploit the rich
787 information available in multimodal datasets in order to align brain images
788 with maximal accuracy. A truly multimodal tool must be able to leverage in-
789 formation about both magnitude and directionality in each voxel, since we are
790 able to capture and represent both of these with MRI. To this end, MMORF
791 is able to explicitly, and simultaneously, optimise both the displacement and
792 rotational effects of a warp field. Importantly, not *optimising* for rotation
793 (as is the case for any scalar registration method) does not mean that there
794 *is* no rotation, only that it cannot be controlled to either improve directional
795 alignment or prevent the introduction of misalignment. Simultaneously opti-
796 mising over multiple modalities through a single warp both allows the unique
797 information from each modality to influence the resulting transformation (as
798 when each modality is aligned separately), and ensures that all modalities
799 remain co-registered following warping (as when using only a single modality
800 to drive alignment).

801 But even when combining information across modalities, image regis-
802 tration is still a highly underdetermined problem and, therefore, alignment
803 accuracy (as measured by how well anatomy are co-localised) is highly depen-

804 dent on the choice of warp regularisation. We therefore included in MMORF
805 a regularisation model that promotes highly biologically plausible deforma-
806 tions, thereby effectively controlling excessive levels of distortion in both
807 shape and size.

808 The combination of our multimodal approach and our regularisation met-
809 ric is computationally intensive, particularly when it comes to calculating the
810 Hessian matrix required by MMORF’s Gauss-Newton optimisation strategy.
811 We address this through GPU parallelisation of the most computationally
812 intensive parts of the algorithm, allowing MMORF to run to 1 mm warp res-
813 olutions within 5 to 45 minutes (depending on image resolution and number
814 of modalities).

815 We have evaluated MMORF across four domains—T1w-label overlap,
816 DTI similarity, tfMRI cluster mass and image distortion. Performance was
817 benchmarked against three established registration tools—FNIRT, ANTs,
818 and DR-TAMAS. These tools represent MMORF’s predecessor, a state-
819 of-the-art unimodal method, and the most similar multimodal alternative,
820 respectively. A common theme was that methods performed well when
821 tested on metrics derived from modalities used in their respective cost func-
822 tions—e.g., FNIRT performed well on T1w-label overlap, and DR-TAMAS
823 performed well on DTI similarity. MMORF performed best or near-best in
824 both scalar and tensor evaluations, best in the held-out tfMRI evaluation,
825 and best in terms of distortion.

826 FNIRT was able to slightly outperform MMORF in cortical T1w-label
827 overlap, but was outperformed by MMORF across all other metrics. FNIRT
828 induces the most distortion out of the methods tested. This likely contributed
829 to the relatively poor DTI similarity results, since excessive deformations are
830 likely to cause incorrect rotation of the tensors. Interestingly, having the
831 best cortical label overlap has not translated to the best tfMRI performance,
832 with FNIRT actually showing the poorest performance, despite both domains
833 being evaluated on the cortex. The strong cortical T1w-label performance
834 is, therefore, likely due to overfitting of the T1w image similarity metric, and
835 highlights the importance of holistically evaluating registration performance,
836 including the use of a held-out modality. These results demonstrate the
837 significant improvement in performance of MMORF over its predecessor.
838 Since both the inputs and outputs to MMORF are fully compatible with
839 FNIRT and FSL, the benefits of this improvement can be realised by simply
840 substituting in the new method.

841 ANTs performed near-identically to MMORF in both cortical and sub-

842 cortical T1w-label overlap measures, making them the best performing meth-
843 ods in this domain. The inclusion of DTI data in MMORF registration has,
844 therefore, not compromised the alignment of the T1w channel, which is not
845 necessarily a given (see the comparison to DR-TAMAS below). The ben-
846 efits of including the DTI data are evident in the DTI similarity results,
847 where MMORF clearly has the advantage over ANTs. MMORF consistently
848 outperforms ANTs in the tfMRI evaluation, indicating better anatomical
849 consistency in grey matter, despite similar label overlap performance. Lev-
850 els of distortion are very comparable between these two methods, despite
851 there being more information to drive white matter deformations harder in
852 MMORF.

853 DR-TAMAS was the poorest performing nonlinear method in terms of
854 cortical T1w-label overlap, and second-poorest subcortically. Since it largely
855 shares its scalar registration algorithm with ANTs, this suggests that this is
856 due to the inclusion of DTI data in driving the registration. As has already
857 been noted, this was not the case for MMORF. This could therefore be at-
858 tributed to the differences in DTI cost functions employed by DR-TAMAS
859 and MMORF—where MMORF uses the whole tensor, while DR-TAMAS
860 uses a combination of the mean diffusivity (a scalar) and the deviatoric ten-
861 sor (shape and direction information only). DR-TAMAS came a close second
862 to MMORF across all DTI similarity metrics, with performance clearly better
863 than the scalar-only methods. DR-TAMAS was also the closest performing
864 method to MMORF in the tfMRI evaluation, demonstrating again that corti-
865 cal segmentation performance does not necessarily translate to well matched
866 cortical activations. In terms of amount of deformation, DR-TAMAS induced
867 more deformations than MMORF, despite having been outperformed across
868 all accuracy measures.

869 One limitation of this and, in fact, most evaluations of registration meth-
870 ods is that we are likely to be able to better tune our own method than those
871 against which we are comparing. However, starting from the default settings,
872 we have made an effort to optimise the performance of each, and kept any
873 adjustments that proved beneficial. As such, we believe that our results are
874 representative of what can be expected in real-world use. While there was
875 always at least one tool with similar performance to MMORF in each test,
876 none were consistently as high-performing across the board.

877 Note that, although our primary objective was human brain alignment,
878 MMORF does not rely on any human brain priors (*e.g.*, tissue maps or
879 assumed brain size), and may be applied to any domain of medical imag-

880 ing. For example, MMORF has been successfully used in the generation of
881 multimodal, non-human primate brain templates of several phylogenetically
882 distant species for the purpose of performing comparative anatomy of white
883 matter tracts (Roumazeilles et al., 2022). Similarly, while the HCP enables
884 us to compare the relative performance of each registration tools across mul-
885 tiple imaging modalities using a single dataset, our evaluation here is largely
886 confined to the cerebrum as there is little to no labelled data in the cere-
887 bellum, brain stem, and spinal cord. MMORF may, in fact, be particularly
888 beneficial for studying these non-cerebral regions, and we hope to explore
889 this is in future.

890 In this work we have only considered MMORF’s performance in compari-
891 son to other volumetric registration tools. There is strong evidence that using
892 surface-based registration has a beneficial effect on cortical alignment accu-
893 racy (Coalson et al., 2018), and we are not advocating the use of MMORF
894 instead of those methods when performing cortical or surface-based cross-
895 subject fMRI analyses, for example. However, volumetric registration is, as
896 yet, still the only way to align non-cortical brain regions, and is likely to
897 remain just as valuable as surface registration in neuroimaging research for
898 the foreseeable future, especially when data are registered intra-individually
899 and when surface reconstructions fail, for example, due to the presence of
900 intra-axial lesions (such as brain tumours).

901 Given MMORF’s performance, reasonable execution times, and simple
902 compatibility with FSL, we believe it to be an excellent choice of tool for vol-
903 umetric registration—regardless of the domain of any follow-on analysis. As
904 such, it is easy to recommend MMORF to all users working with multimodal
905 neuroimaging data, and in particular those who already rely on FSL for their
906 analyses. They will benefit from state-of-the-art registration accuracy across
907 all domains of their downstream analyses, with very little modification to
908 any existing pipelines. MMORF is available as part of FSL in releases newer
909 than 6.0.7.

910 6. Data and Code Availability

911 MMORF binaries and the associated source code are available for down-
912 load via the standard FSL installer³. A stand-alone Singularity image, in-
913 structions for running MMORF, and example configuration files are available

³<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FslInstallation>

914 on the FMRIB GitLab server⁴. All analysis scripts, figures, figure source code
915 and latex source code that were used to generate this manuscript are available
916 on the FMRIB GitLab server⁵.

917 7. Author Contributions

918 **Frederik Lange:** Conceptualisation, Methodology, Software, Validation,
919 Formal analysis, Investigation, Writing (original draft, review and editing),
920 Visualisation. **Christoph Arthofer:** Methodology, Software, Validation,
921 Writing (review and editing). **Andreas Bartsch and Gwenaëlle Douaud:**
922 Validation, Writing (review and editing). **Paul McCarthy:** Software, Data
923 Curation. **Stephen Smith:** Conceptualisation, Methodology, Validation,
924 Resources, Writing (review and editing), Supervision, Project administra-
925 tion, Funding acquisition. **Jesper Andersson:** Conceptualisation, Method-
926 ology, Software, Validation, Formal analysis, Writing (review and editing),
927 Supervision, Project administration.

928 8. Acknowledgements

929 This research was funded by a Wellcome Trust Collaborative Award
930 (215573/Z/19/Z). The Wellcome Centre for Integrative Neuroimaging is sup-
931 ported by core funding from the Wellcome Trust (203139/Z/16/Z). Compu-
932 tation used the Oxford Biomedical Research Computing (BMRC) facility, a
933 joint development between the Wellcome Centre for Human Genetics and the
934 Big Data Institute supported by Health Data Research UK and the NIHR
935 Oxford Biomedical Research Centre. The views expressed are those of the
936 author(s) and not necessarily those of the NHS, the NIHR or the Department
937 of Health. Data were provided in part by the Human Connectome Project,
938 WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil
939 Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that
940 support the NIH Blueprint for Neuroscience Research; and by the McDonnell
941 Center for Systems Neuroscience at Washington University.

942 9. Declaration of Competing Interests

943 We have no competing interests to declare.

⁴<https://git.fmrib.ox.ac.uk/flange/mmorph.beta>

⁵<https://git.fmrib.ox.ac.uk/flange/mmorph.toolbox>

944 References

945 Alexander, A.L., Hasan, K., Kindlmann, G., Parker, D.L., Tsuruda, J.S.,
946 2000. A geometric analysis of diffusion tensor measurements of the human
947 brain. *Magnetic Resonance in Medicine* 44, 283–291.

948 Alexander, D.C., Pierpaoli, C., Basser, P.J., Gee, J.C., 2001. Spatial transfor-
949 mations of diffusion tensor magnetic resonance images. *IEEE Transactions*
950 on Medical Imaging 20, 1131–1139.

951 Allen, M., Poggiali, D., Whitaker, K., Marshall, T.R., Van Langen, J., Kievit,
952 R.A., 2021. Raincloud plots: A multi-platform tool for robust data visu-
953 alization. *Wellcome Open Research* 4, 63.

954 Andersson, J.L.R., Jenkinson, M., Smith, S.M., 2007. Non-Linear Registra-
955 tion Aka Spatial Normalisation. Technical Report June.

956 Andersson, J.L.R., Jenkinson, M., Smith, S.M., 2019. High resolution non-
957 linear registration with simultaneous modelling of intensities. *bioRxiv* ,
958 646802.

959 Andersson, J.L.R., Skare, S., Ashburner, J., 2003. How to correct suscepti-
960 bility distortions in spin-echo echo-planar images: Application to diffusion
961 tensor imaging. *NeuroImage* 20, 870–888.

962 Andersson, J.L.R., Sotropoulos, S.N., 2016. An integrated approach to cor-
963 rection for off-resonance effects and subject movement in diffusion MR
964 imaging. *NeuroImage* 125, 1063–1078.

965 Arthofer, C., Smith, S.M., Andersson, J.L.R., Lange, F.J., 2021. Multimodal
966 MRI template construction from UK Biobank: Oxford-MM-0, in: 27th
967 Annual Meeting of The Organization for Human Brain Mapping.

968 Arthofer, C., Smith, S.M., Douaud, G., Bartsch, A.J., Andersson, J.L.R.,
969 Lange, F.J., 2022. Age-dependent multimodal MRI template construc-
970 tion from UK Biobank, in: 28th Annual Meeting of The Organization for
971 Human Brain Mapping.

972 Ashburner, J., 2007. A fast diffeomorphic image registration algorithm. *Neu-
973 roImage* 38, 95–113.

974 Ashburner, J., Andersson, J.L.R., Friston, K.J., 1999. High-dimensional
975 image registration using symmetric priors. *NeuroImage* 9, 619–628.

976 Ashburner, J., Andersson, J.L.R., Friston, K.J., 2000. Image registration
977 using a symmetric prior - In three dimensions. *Human Brain Mapping* 9,
978 212–225.

979 Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C., 2008. Symmetric
980 diffeomorphic image registration with cross-correlation: Evaluating auto-
981 mated labeling of elderly and neurodegenerative brain. *Medical Image
982 Analysis* 12, 26–41.

983 Bajcsy, R., Kovačič, S., 1989. Multiresolution elastic matching. *Computer
984 Vision, Graphics, and Image Processing* 46, 1–21.

985 Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L.,
986 Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D.,
987 Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S.,
988 Johansen-Berg, H., Snyder, A.Z., Van Essen, D.C., 2013. Function in
989 the human connectome: Task-fMRI and individual differences in behavior.
990 *NeuroImage* 80, 169–189.

991 Beg, M.F., Miller, M., Trouvé, A., Younes, L., 2005. Computing large defor-
992 mation metric mappings via geodesic flows of diffeomorphisms. *Inter-
993 national Journal of Computer Vision* 61, 139–157.

994 Bookstein, F.L., 1997. Quadratic variation of deformations, in: *Lecture
995 Notes in Computer Science (Including Subseries Lecture Notes in Artificial
996 Intelligence and Lecture Notes in Bioinformatics)*. volume 1230, pp. 15–28.

997 Bro-Nielsen, M., Gramkow, C., 1996. Fast fluid registration of medical im-
998 ages, in: *Lecture Notes in Computer Science (Including Subseries Lecture
999 Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. vol-
1000 ume 1131, pp. 267–276.

1001 Christensen, G.E., 1999. Consistent linear-elastic transformations for image
1002 matching, in: Kuba, A., Šámal, M., Todd-Pokropek, A. (Eds.), *Lecture
1003 Notes in Computer Science (Including Subseries Lecture Notes in Arti-
1004 ficial Intelligence and Lecture Notes in Bioinformatics)*, Springer Berlin
1005 Heidelberg, Berlin, Heidelberg. pp. 224–237.

1006 Chun, I.Y., Fessler, J.A., 2018. Convolutional Dictionary Learning: Acceleration and Convergence. *IEEE Transactions on Image Processing* 27,
1007 1697–1712.

1009 Coalson, T.S., Van Essen, D.C., Glasser, M.F., 2018. The impact of traditional
1010 neuroimaging methods on the spatial localization of cortical areas. *Proceedings of the National Academy of Sciences* 115, E6356–E6365.

1012 Destrieux, C., Fischl, B., Dale, A., Halgren, E., 2010. Automatic parcellation
1013 of human cortical gyri and sulci using standard anatomical nomenclature.
1014 *NeuroImage* 53, 1–15.

1015 Elam, J.S., Glasser, M.F., Harms, M.P., Sotiroopoulos, S.N., Andersson,
1016 J.L.R., Burgess, G.C., Curtiss, S.W., Oostenveld, R., Larson-Prior, L.J.,
1017 Schoffelen, J.M., Hodge, M.R., Cler, E.A., Marcus, D.M., Barch, D.M.,
1018 Yacoub, E., Smith, S.M., Ugurbil, K., Van Essen, D.C., 2021. The Human
1019 Connectome Project: A retrospective. *NeuroImage* 244, 118543.

1020 Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C.,
1021 van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo,
1022 A., Makris, N., Rosen, B., Dale, A.M., 2002. Whole Brain Segmentation:
1023 Automated Labeling of Neuroanatomical Structures in the Human Brain.
1024 *Neuron* 33, 341–355.

1025 Fischl, B., Salat, D.H., van der Kouwe, A.J.W., Makris, N., Ségonne, F.,
1026 Quinn, B.T., Dale, A.M., 2004. Sequence-independent segmentation of
1027 magnetic resonance images. *NeuroImage* 23, S69–S84.

1028 Glasser, M.F., Sotiroopoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson,
1029 J.L.R., Xu, J., Jbabdi, S., Webster, M.A., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., 2013. The minimal preprocessing pipelines for
1030 the Human Connectome Project. *NeuroImage* 80, 105–124.

1032 Hunter, D.R., Lange, K., 2004. A Tutorial on MM Algorithms. *American
1033 Statistician* 58, 30–37.

1034 Irfanoglu, M.O., Nayak, A., Jenkins, J., Hutchinson, E.B., Sadeghi, N.,
1035 Thomas, C.P., Pierpaoli, C., 2016. DR-TAMAS: Diffeomorphic Registration
1036 for Tensor Accurate Alignment of Anatomical Structures. *NeuroImage*
1037 132, 439–454.

1038 Jenkinson, M., Bannister, P.R., Brady, J.M., Smith, S.M., 2002. Improved
1039 optimization for the robust and accurate linear registration and motion
1040 correction of brain images. *NeuroImage* 17, 825–841. [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).

1041 Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith,
1042 S.M., 2012. FSL. *NeuroImage* 62, 782–790.

1043 Jenkinson, M., Smith, S.M., 2001. A global optimisation method for robust
1044 affine registration of brain images. *Medical Image Analysis* 5, 143–156.

1045 Klein, A., Andersson, J.L.R., Ardekani, B.A., Ashburner, J., Avants, B.B.,
1046 Chiang, M.C., Christensen, G.E., Collins, D.L., Gee, J.C., Hellier, P.,
1047 Song, J.H., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P.M.,
1048 Vercauteren, T., Woods, R.P., Mann, J.J., Parsey, R.V., 2009. Evaluation
1049 of 14 nonlinear deformation algorithms applied to human brain MRI
1050 registration. *NeuroImage* 46, 786–802.

1051 Lange, F.J., Ashburner, J., Smith, S.M., Andersson, J.L.R., 2020. A Sym-
1052 metric Prior for the Regularisation of Elastic Deformations: Improved
1053 anatomical plausibility in nonlinear image registration. *NeuroImage* 219,
1054 116962.

1055 Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub,
1056 E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiroopoulos, S.N., Andersson, J.L.R.,
1057 Griffanti, L., Douaud, G., Okell, T.W., Weale, P., Dragonu, I., Garratt,
1058 S., Hudson, S., Collins, R., Jenkinson, M., Matthews, P.M., Smith, S.M.,
1059 2016. Multimodal population brain imaging in the UK Biobank prospective
1060 epidemiological study. *Nature Neuroscience* 19, 1523–1536.

1061 Miller, M., Banerjee, A., Christensen, G., Joshi, S., Khaneja, N., Grenan-
1062 der, U., Matejic, L., 1997. Statistical methods in computational anatomy.
1063 *Statistical Methods in Medical Research* 6, 267–299.

1064 Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes,
1065 D.J., Fox, N.C., Ourselin, S., 2010. Fast free-form deformation using graph-
1066 ics processing units. *Computer Methods and Programs in Biomedicine* 98,
1067 278–284.

1068 Møller, M.F., 1993. A scaled conjugate gradient algorithm for fast supervised
1069 learning. *Neural Networks* 6, 525–533.

1070 Nocedal, J., Wright, S.J., 2006. Numerical Optimization. 9781447122234,
1071 Springer Science+Business Media, USA, New York.

1072 NVIDIA, 2019. CUDA C Programming Guide .

1073 Robinson, E.C., Garcia, K., Glasser, M.F., Chen, Z., Coalson, T.S.,
1074 Makropoulos, A., Bozek, J., Wright, R., Schuh, A., Webster, M., Hutt-
1075 ter, J., Price, A., Cordero Grande, L., Hughes, E., Tusor, N., Bayly, P.V.,
1076 Van Essen, D.C., Smith, S.M., Edwards, A.D., Hajnal, J., Jenkinson, M.,
1077 Glocker, B., Rueckert, D., 2018. Multimodal surface matching with higher-
1078 order smoothness constraints. NeuroImage 167, 453–465.

1079 Rohlfing, T., 2012. Image Similarity and Tissue Overlaps as Surrogates for
1080 Image Registration Accuracy: Widely Used but Unreliable. IEEE Trans-
1081 actions on Medical Imaging 31, 153–163.

1082 Rohlfing, T., Maurer, C.R., Bluemke, D.A., Jacobs, M.A., 2003. Volume-
1083 Preserving Nonrigid Registration of MR Breast Images Using Free-Form
1084 Deformation With an Incompressibility Constraint. IEEE Transactions on
1085 Medical Imaging 22, 730–741.

1086 Roumazeilles, L., Lange, F.J., Benn, R.A., Andersson, J.L.R., Bertelsen,
1087 Manger, P.R., Flach, E., Khrapitchev, A.A., Bryant, K.L., Sallet,
1088 J., Mars, R.B., 2022. Cortical Morphology and White Matter Tractogra-
1089 phy of Three Phylogenetically Distant Primates: Evidence for a Simian
1090 Elaboration. Cerebral Cortex 32, 1608–1624.

1091 Smith, W., Wormald, N., 1998. Geometric separator theorems and applica-
1092 tions, in: Proceedings 39th Annual Symposium on Foundations of Com-
1093 puter Science (Cat. No.98CB36280), IEEE Comput. Soc. pp. 232–243.

1094 Szeliski, R., Coughlan, J., 1997. Spline-Based Image Registration. Interna-
1095 tional Journal of Computer Vision 22, 199–218.

1096 Tagare, H.D., Girois, D., Skrinjar, O., 2009. Symmetric Non-rigid Regis-
1097 tration: A Geometric Theory and Some Numerical Techniques. Journal of
1098 Mathematical Imaging and Vision 34, 61–88.

1099 Unser, M., Aldroubi, A., Eden, M., 1993a. B-spline signal processing. I.
1100 Theory. IEEE Transactions on Signal Processing 41, 821–833.

1101 Unser, M., Aldroubi, A., Eden, M., 1993b. B-spline signal processing. II.
1102 Efficiency design and applications. *IEEE Transactions on Signal Processing*
1103 41, 834–848.

1104 Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E.,
1105 Ugurbil, K., 2013. The WU-Minn Human Connectome Project: An
1106 overview. *NeuroImage* 80, 62–79.

1107 Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E., Bu-
1108 cholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S., Della Penna,
1109 Feinberg, D., Glasser, M.F., Harel, N., Heath, A., Larson-Prior, L.,
1110 Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S.E.,
1111 Prior, F., Schlaggar, B.L., Smith, S.M., Snyder, A., Xu, J., Yacoub, E.,
1112 2012. The Human Connectome Project: A data acquisition perspective.
1113 *NeuroImage* 62, 2222–2231.

1114 Vovk, U., Pernuš, F., Likar, B., 2007. A Review of Methods for Correction of
1115 Intensity Inhomogeneity in MRI. *IEEE Transactions on Medical Imaging*
1116 26, 405–421.

1117 Winkler, A.M., Ridgway, G.R., Webster, M.A., Smith, S.M., Nichols, T.E.,
1118 2014. Permutation inference for the general linear model. *NeuroImage* 92,
1119 381–397.

1120 Yeo, B.T., Vercauteren, T., Fillard, P., Peyrat, J.M., Pennec, X., Golland,
1121 P., Ayache, N., Clatz, O., 2009. DT-REFinD: Diffusion Tensor Registration
1122 With Exact Finite-Strain Differential. *IEEE Transactions on Medical
1123 Imaging* 28, 1914–1928.

1124 Zhang, H., Yushkevich, P.A., Alexander, D.C., Gee, J.C., 2006. Deformable
1125 registration of diffusion tensor MR images with explicit orientation opti-
1126 mization. *Medical Image Analysis* 10, 764–785.