Next-Gen Tools

Using synthetic MR images for distortion correction

David F. Montez [a,b,*,1], Andrew N. Van [a,d,1], Ryland L. Miller [a,b], Nicole A. Seider [a], Scott Marek [b], Annie Zheng [a], Dillan J. Newbold [a,c], Kristen Scheidter [a,b], Eric Feczko [e,f], Anders J. Perrone [e,g], Oscar Miranda-Dominguez [e,f], Eric A. Earl [e,g], Benjamin P. Kay [a], Abhinav K. Jha [d,h], Aristeidis Sotiras [h,i], Timothy O. Laumann [b], Deanna J. Greene [j], Evan M. Gordon [h], M. Dylan Tisdall [k], Andre van der Kouwe [l,m], Damien A. Fair [e,f,n], Nico U.F. Dosenbach [a,d,h,o]

[a] Department of Neurology, Washington University School of Medicine, St. Louis, MO 63110, United States of America
[b] Department of Psychiatry, Washington University School of Medicine, St. Louis, MO 63110, United States of America
[c] Department of Neurology, New York University Langone Medical Center, New York, NY 10016, United States of America
[d] Department of Biomedical Engineering, Washington University School of Medicine, St. Louis, MO 63110, United States of America
[e] Masonic Institute for the Developing Brain, University of Minnesota Medical School, Minneapolis, MN 55455, United States of America
[f] Department of Pediatrics, University of Minnesota Medical School, Minneapolis, MN 55455, United States of America
[g] Department of Psychiatry, Oregon Health and Science University, Portland, OR 97239, United States of America
[h] Department of Radiology, Washington University School of Medicine, St. Louis, MO 63110, United States of America
[i] Institute for Informatics, Washington University School of Medicine, St. Louis, MO 63110, United States of America
[j] Department of Cognitive Science, University of California, San Diego, La Jolla CA 92093, United States of America
[k] Department of Radiology, Perelman School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, United States of America
[l] Department of Radiology, Massachusetts General Hospital, Athinoula A. Martinos Center for Biomedical Imaging, Charlestown, MA 02129, United States of America
[m] Department of Radiology, Harvard Medical School, Boston, MA 02115, United States of America
[n] Institute of Child Development, University of Minnesota Medical School, Minneapolis, MN 55455, United States of America
[o] Department of Pediatrics, Washington University School of Medicine, St. Louis, MO 63110, United States of America

ARTICLE INFO

ABSTRACT

Functional MRI (fMRI) data acquired using echo-planar imaging (EPI) are highly distorted by magnetic field inhomogeneities. Distortion and differences in image contrast between EPI and T1-weighted and T2-weighted (T1w/T2w) images makes their alignment a challenge. Typically, field map data are used to correct EPI distortions. Alignments achieved with field maps can vary greatly and depends on the quality of field map data. However, many public datasets lack field map data entirely. Additionally, reliable field map data is often difficult to acquire in high-motion pediatric or developmental cohorts. To address this, we developed *Synth*, a software package for distortion correction and cross-modal image registration that does not require field map data. *Synth* combines information from T1w and T2w anatomical images to construct an idealized undistorted synthetic image with similar contrast properties to EPI data. This synthetic image acts as an effective reference for individual-specific distortion correction. Using pediatric (ABCD: Adolescent Brain Cognitive Development) and adult (MSC: Midnight Scan Club; HCP: Human Connectome Project) data, we demonstrate that *Synth* performs comparably to field map distortion correction approaches, and often outperforms them. Field map-less distortion correction with *Synth* allows accurate and precise registration of fMRI data with missing or corrupted field map information.

1. Introduction

BOLD-weighted (blood-oxygenation level dependent) functional MRI (fMRI) data obtained using echo planar imaging (EPI) is severely distorted by inhomogeneities affecting the primary magnetic field (Andersson and Skare, 2002; Andersson et al., 2003; Callaghan, 1990; Jezzard and Balaban, 1995; Jezzard et al., 1998). EPI distortion — which consists of localized spatial deformation, loss of BOLD signal

intensity, and signal 'pile-up'— prominently affects areas containing large local differences in magnetic susceptibility. Due to the apposition of diamagnetic tissue and paramagnetic air in the sinuses and ear canals, regions of the image containing the orbitofrontal cortex and the inferior temporal lobes often suffer the most severe distortions. The susceptibility-induced artifacts are spatially non-uniform and therefore interfere with the performance of registration algorithms used during fMRI preprocessing to bring BOLD EPI images into alignment with their associated anatomical images (T1w, T2w).

Establishing the correspondence between brain anatomy and function is an important component of interpreting neuroimaging findings. Regardless of study design or choice of analysis space, registration and removal of EPI distortion and registration to T1w/T2w anatomical images are crucial steps in the analysis of fMRI data. For example, poor alignment counteracts the benefits of analyses performed with reference to participant-specific anatomy. Improper alignment of EPI and anatomical images also degrades the performance of procedures that project volumetric fMRI data onto mesh surfaces derived from tissue segmentation (Dickie et al., 2019). The effects of poor registration and distortion correction also carry forward into group analyses. In group studies, anatomical images from many participants are separately aligned to a reference atlas. However, suboptimal alignment between a participant's EPI and anatomical images will propagate as nuisance variability that negatively affects group-level statistics for both task and resting state analyses (Cusack et al., 2003; Togo et al., 2017).

Significant effort has been devoted to developing methods to correct distortions affecting fMRI data. Currently, two primary methods are commonly used, both of which involve acquiring field map data at scan time. The first method involves acquiring a pair of EPI images with opposing phase encoding directions. Because the largest EPI distortions occur in the phase encoding direction, reversing the phase encoding direction also reverses the direction of the EPI distortions. By combining information from both directions, a displacement field that corrects for the underlying EPI distortion can be constructed (Andersson et al., 2003). The second field map method relies on the linear relationship between the phase of gradient echo data, local magnetic field inhomogeneities, and echo time. By recording MR data with two different echo times, a displacement field can be constructed that corrects for distortion caused by inhomogeneity (Jezzard and Balaban, 1995). The end product of each of these approaches is a nonlinear warp that indicates how each voxel must be displaced in order to correct for the image distortion.

Without mitigation, the validity and quality of corrections produced by either of the standard field map methods hinges on several practical points: First, subject movement during the acquisition of the field map data or during the time between the acquisition of the field map data and the corresponding EPI data can negatively affect corrections in two ways: (1) motion can reduce the quality of the field map data itself by introducing artifacts which reduce the accuracy of corrections; and (2) movement may also change the spatial structure of the distortion so that the geometry of the brain images differs in the field map data and the EPI data (Andersson and Sotiropoulos, 2016). This may reduce the accuracy of the corrections simply by reducing the accuracy of the alignments between EPI data and field map data meant to provide the correction. Second, research has shown that dual-echo and phase-reversal methods for estimating EPI distortion perform differently across brain regions and that the method used to generate the echoes (spin-echo vs. gradient-echo) used in field map acquisitions can affect the accuracy of distortion correction (Schallmo et al., 2021). Third, extreme levels of distortion can compress features so excessively – so-called signal "pile-up"– that they cannot be correctly unwarped. Finally, the process of estimating the distortion correction from field map data, in most cases, relies on internal regularization parameters of the algorithm used to estimate the distortion. These parameters affect the overall smoothness of the final correcting warp, and may not be optimal for a particular set of image acquisition parameters or level of image detail.

The need to acquire field maps introduces a crucial failure point during data acquisition. Failure to acquire valid field map data can disqualify an entire dataset from inclusion in an analysis. This problem is further amplified in study designs requiring multiple field maps be collected across multiple sessions or even within sessions. In these cases, each field map represents a potential failure point where poor data collection can affect the quality of post-acquisition analysis. For instance, within the ABCD Annual Release 2.0 dataset (DOI 10.15154/1503209) 341 participants were missing valid field map data. At the time of writing, of the top five most downloaded fMRI datasets available on OpenNeuro.org (Flanker task (event-related), UCLA Consortium for Neuropsychiatric Phenomics LA5c Study, Classification learning, Forrest Gump, Multisubject, multimodal face processing), only one includes field map data. An effective implementation of a field map-less approach is highly desirable because it would allow for the correction of images for which field map information was either missing, corrupted, or never collected.

Motivated by the limitations of field maps and the potential to reinvestigate datasets lacking field maps, fMRI researchers have explored direct mapping approaches in which distortion is corrected by non-linearly registering a participant's distorted EPI images directly to their undistorted anatomical image. Early implementations of direct mapping demonstrated that the approach could reduce distortion and improve global measures of image similarity (e.g., mutual information or squared-error) between EPI and anatomical images (Gholipour et al., 2008a,b; Kybic et al., 2000; Studholme et al., 2000).

More recent implementations of the direct mapping approach include external group average field map data as a constraint on possible solutions (e.g., fieldmap-less SyN-Susceptibility Distortion Correction (SDC)). Quality assessment of distortion corrections produced by direct mapping demonstrate that final image alignment quality can be unreliable and tends to perform more poorly than high quality corrections using a field map (Chambers et al., 2015; Hong et al., 2015; Huntenburg, 2014; Wang et al., 2017). One likely cause for the variable performance of this approach is the fact that T1w, T2w, and EPI images look quite different from one another. Because the average signal intensities associated with the tissues and fluids comprising the brain vary considerably across acquisition parameters, it is difficult to construct registration cost functions that accurately reflect the true error introduced by EPI distortion (West et al., 1996). This problem is compounded by the low resolution of EPI images and signal pileup or dropout in highly distorted areas and challenges in selecting appropriate spatial regularization terms for nonlinear optimization.

Image correction using undistorted synthetic image references is another more recent approach to field map-less distortion correction. This strategy entails the estimation of an undistorted auxiliary EPI image that serves as a reference for unwarping. The SynB0-DisCo algorithm is a notable example of this approach that employs trained deep learning neural networks to transform undistorted T1w images into a synthetic EPI image which is used as an input image for FSL's *topup* algorithm (Schilling et al., 2020). Initial assessments of this approach suggest that it can potentially produce distortion corrections that are comparable to those produced by high quality field map data. However, because T1w and EPI image contrast can vary significantly across acquisitions and depends on the accuracy of the intensity bias correction applied to the images, it is unclear how well these trained deep learning approaches will perform on arbitrary datasets. In addition, there is potential for synthetic images produced by the deep learning networks to include spurious image artifacts which may affect the reliability of distortion corrections in practice (Antun et al., 2020; Bhadra et al., 2021).

We sought to combine direct mapping and synthetic reference image approaches to produce a more reliable field map-less distortion correction algorithm — one that does not rely on priors established by

training data. To accomplish this, we developed a modeling framework that allows us to combine information from a participant's T1w, T2w, and EPI images in order to construct a synthetic image that has the contrast properties of a EPI image and the undistorted geometry and high resolution (~1 mm³) of typical T1w/T2w images. We hypothesized that these synthetic EPI images would serve as ideal targets for estimating field map corrections with currently available nonlinear warping software, and thereby improve direct mapping quality. Synthetic reference images created in this way do not rely on trained models. Consequently, this approach flexibly adapts to the particular contrast properties of a given dataset and may be less prone to the influence of deep learning network artifacts.

Here, we detail *Synth*, an implementation of our synthetic reference image approach to correcting EPI distortion. We begin by describing the mathematical framework we use to generate synthetic EPI images as well as an approach for using them to correct distortion. Then, we demonstrate that undistorted synthetic images constructed using this framework are quantitatively more similar to real EPI images than T1w or T2w images. Building on these results, we compare this approach to distortion correction against other commonly used MRI field map-based and field map-less correction procedures. We performed these comparisons in a subset of the Adolescent Brain Cognitive Development (ABCD) study dataset (Jernigan et al., 2018) in order to assess performance across a variety of scanners, sites and brain geometries (Marek et al., 2019). Additionally, we compared performance in the Midnight Scan Club (MSC) dataset, which consists of ten highly sampled individuals (Gordon et al., 2017). The repeated sampling of MSC participants allowed us to assess the reliability of various methods of EPI distortion correction when applied to the same participants across multiple acquisitions. Lastly, we compare the *Synth* distortion correction to standard field map corrections in a ten subject sample taken from the Human Connectome data set. We demonstrate in each case that *Synth* software can be used to correct fMRI distortions using a synthetic image as an alignment target. The *Synth* software may be used to augment existing fMRI preprocessing pipelines, or explored by researchers interested in incorporating variations on these themes into their MRI registration procedures.

2. Methods

2.1. Creating an undistorted synthetic EPI image using Synth

Our approach to correcting EPI distortion is to create a synthetic EPI image based on information combined from a participant's undistorted T1w and T2w images and then use this synthetic image as a reference for nonlinearly aligning a participants' real EPI image. Here, we refer to any real image that *Synth* attempts to synthesize as a target image. The ideal synthetic EPI image will have three properties: (1) it will match the real EPI target image in terms of overall signal intensity; (2) it will exhibit similar contrast between fluids and tissue types that closely correspond to what is observed in the EPI target image and; (3) It will account for the difference in spatial resolution between typically high-resolution anatomical images and typically low-resolution EPI images. Thus, the synthetic image can be represented by the following model:

$$\mathbf{s} = \mathbf{B}\left(\mathbf{EF}\theta; \alpha, \beta\right) \tag{1}$$

The left hand side of Eq. (1), $\mathbf{s}$, represents a final synthetic image. Working from inside to outside, the right hand side of Eq. (1) represents the following: the matrix, $\mathbf{F}$, is the decomposition of pre-aligned T1w and T2w anatomical images into a set of basis vectors (e.g., radial basis functions — as in Fig. 1a, or b-splines, etc.) in order to model the continuous relationship between the voxel intensities of the anatomical images and EPI target image; θ comprises the weights for each column of $\mathbf{F}$. $\mathbf{E}$ is a blurring operator modeling the effective difference in spatial resolution between the anatomical images and EPI target image. In our implementation of *Synth*, the blurring operator, $\mathbf{E}$, is effected with an

Epanachnikov smoothing kernel, chosen for its optimal noise-resolution properties (Gureyev et al., 2020). Finally, $\mathbf{B}$ is a global image contrast operator modeled as a cumulative beta distribution, parameterized by scalars α and β. When θ, α, and β are at the optimum solution, $\mathbf{s}$ represents a synthetic image with the geometry of the undistorted source images and the contrast properties of the EPI target image. Solving for the optimal θ, α, and β requires solving a joint optimization problem, which is discussed in the next section.

The columns of the matrix, $\mathbf{F}$, are composed of a radial basis function (RBF) decomposition of the source T1w and T2w anatomical images. This decomposition is visualized in Fig. 1a, where the T1w and T2w input images have been divided into smooth voxel intensity 'bins', each of which corresponds to an RBF component. Each of these components are weighted by the parameter, θ, such that the weighted sum of these components replicates the contrast properties of the EPI target image as visualized in Fig. 1b. *Synth* allows for a user-defined number of intensity 'bins' into which the T1w/T2w images are decomposed. The choice of number of intensity bins is dictated by a tradeoff between computational complexity and the accuracy of the synthetic image. As the number of RBF components in $\mathbf{F}$ increases, so too do the memory requirements. For the synthetic images in the presented results, a 24 component RBF decomposition of both T1w/T2w images (12 components of T1w; 12 components of T2w) at 1 mm isotropic resolution consumed ~20 GB of RAM.

2.2. Estimation of synthetic image and distortion-correction warp

To relate the synthetic EPI image to the real EPI target image, we model the distortion of the synthetic image as a non-linear warp, f, constrained to operate in the phase-encoding direction. This model is described by:

$$\mathbf{y} = f\left(\mathbf{s}; \boldsymbol{\phi}\right) + \boldsymbol{\eta} \tag{2}$$

Here, $\mathbf{y}$ represents the target image (e.g., a distorted EPI image), $\boldsymbol{\phi}$ represents the underlying parameterization of the nonlinear transformation, f, and $\boldsymbol{\eta}$ represents additive gaussian noise. By combining Eqs. (1) and (2), and solving for the parameters that maximize the correlation coefficient between the real and distorted synthetic image, we obtain the following joint optimization problem:

$$\left\{\boldsymbol{\phi}^*, \alpha^*, \beta^*, \theta^*\right\} = \underset{\boldsymbol{\phi}, \alpha, \beta, \theta}{\arg\max}\, \text{corr}\left(\mathbf{y}, f\left(\mathbf{B}\left(\mathbf{EF}\theta; \alpha, \beta\right); \boldsymbol{\phi}\right)\right) \tag{3}$$

To solve Eq. (3), we employ an alternating minimization approach, where one parameter is optimized while the others are held fixed over each iteration. $\boldsymbol{\phi}$ is initialized with a rigid-body transform so that both synthetic and real EPI images are closely registered. Once an updated value for $\boldsymbol{\phi}$ is found through non-linear registration, the inverse of the estimated nonlinear warp, f^{-1}, is then computed. f^{-1} is analogous to a traditional field map correction and maps voxels in a distorted EPI image, $\mathbf{y}$, back to their correct locations. Using the corrected target image, $f^{-1}(\mathbf{y})$, an updated value for θ can be found by solving the linear system $f^{-1}(\mathbf{y}) = \mathbf{EF}\theta$. This updated θ is used to generate an updated intermediary synthetic image. Finally, the parameters, α and β, that control the global contrast for the intermediary synthetic image, are computed by minimizing the least squares difference between the synthetic image and distortion corrected EPI image. The resulting synthetic image is then used for another iteration of non-linear registration.

In regions of the brain with significant distortion and signal dropout or pile-up affecting the EPI image, there exists no valid mapping between a participant's anatomical and EPI images. Therefore *Synth* also allows for the inclusion of a weight volume to reduce the contributions of these areas when estimating the parameters, θ, α and β. For the presented results, we down-weighted the contributions of voxels in high-distortion areas (Supplemental Figure 3).

With the application of f^{-1} to the real EPI data during each iteration, geometric correspondences between anatomical and functional
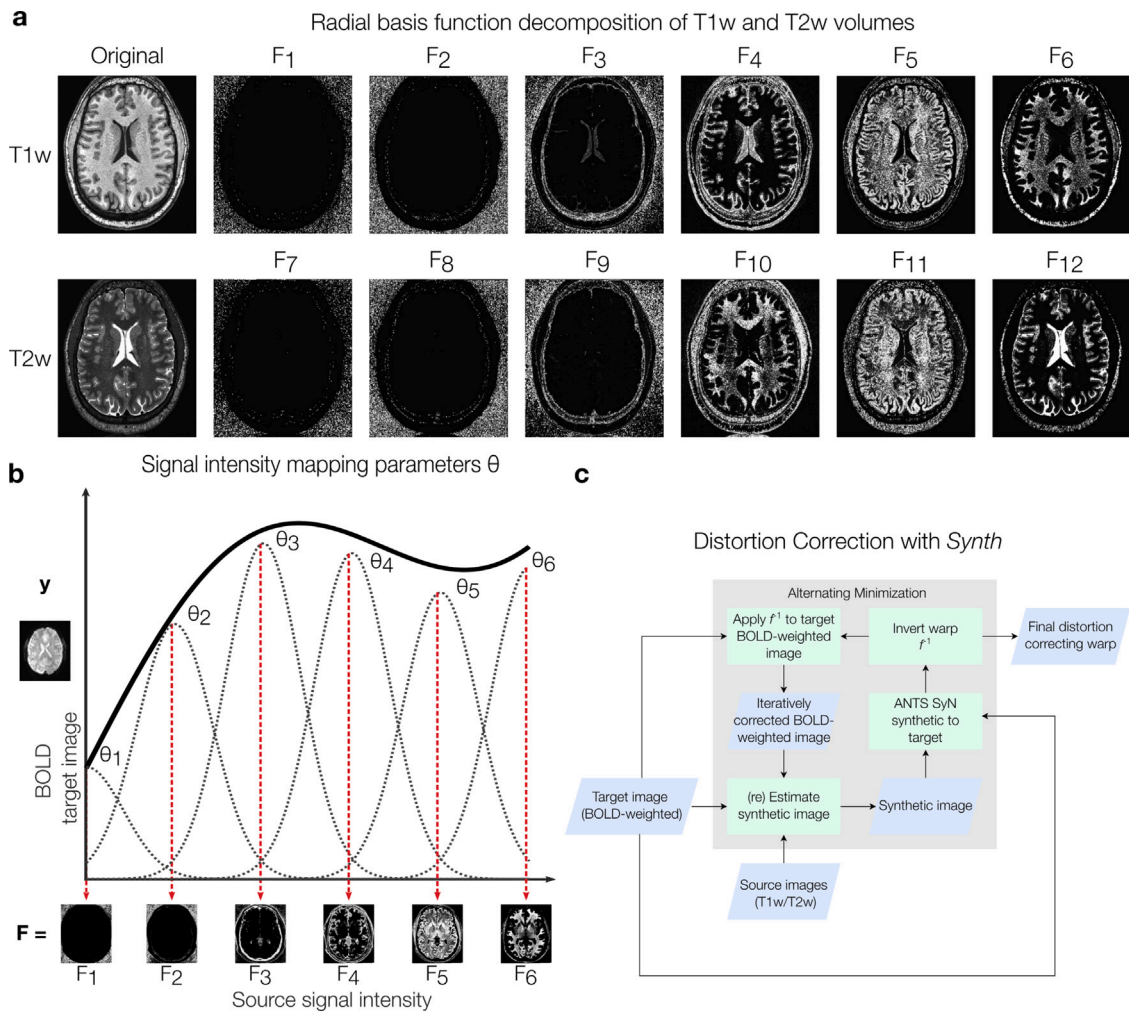
Fig. 1. *Synth* **synthetic image parameters (a)** Illustrative radial basis function (RBF) decomposition of T1w and T2w source images that comprise the columns of **F** (Eq. (1)). RBF decomposition divides source images into smooth 'bins' of signal intensity so that the mapping between source and target image intensities can be estimated. Voxel intensities of a target image, e.g., an EPI image can be modeled as linear combinations of T1w and T2w RBF images. For visualization purposes, we depict a six component decomposition of the T1w image only; *Synth* allows for an arbitrary degree of image decomposition to maximize flexibility in modeling target images with different contrast properties. **(b)** A portion of a hypothetical nonlinear relationship between voxel intensity values observed between source image RBF components (e.g., T2w; x-axis) and voxel intensity values observed in a target image (e.g. EPI). This image depicts the **F**θ portion of Eq. (1). **(c)** Overview visual of the *Synth* algorithm. The displacement field, f, that aligns the initial synthetic image to the target image is estimated using the SyN algorithm. The resulting displacement field, f, is inverted to produce the distortion correcting warp, f^{-1}, which reduces the target EPI image distortion. This improves the correspondence between the target and synthetic images, and by proxy, the source images, allowing for improved estimates of a new synthetic image. The distortion correcting warp, f^{-1}, is updated in an alternating minimization scheme, during which the synthetic image is refined after each improved estimate of f^{-1}.

images are improved. Because the quality and accuracy of the intermediary synthetic EPI images, depends on how closely registered the anatomical and true EPI images are, modest improvements to the synthetic image can be achieved after each iteration of the *Synth* algorithm. By default, *Synth* repeats this process for 3 iterations. The full procedure for solving Eq. (3) is outlined in Algorithm 1 (see Section 6.8 in Supplemental Methods) and Fig. 1c.

In principle, many of the pre-existing nonlinear registration utilities are suitable for estimating f and f^{-1}, owing either to their methods of diffeomorphic warp construction which guarantee the existence of an invertible warp —as is the case with AFNI's 3dQwarp (Zhark the Grotesquely Warped (but still strangely handsome), 2021) and ANTs SyN (Avants et al., 2008)— or their ability to project a potentially non-invertible warp onto a "nearest invertible" space— as is the case with FSL's FNIRT (Paul, 2021). For the results presented here and in the reference preprocessing scripts associated with this manuscript, we used the ANTs SyN algorithm with a local cross correlation metric for estimating all non-linear warps (see Appendix A.7 in Supplemental Methods). ANTs SyN was chosen for its reliable high performance in

non-linear warp estimation (Avants et al., 2011; Klein et al., 2009; Wang et al., 2017).

2.3. Description of MRI datasets and processing

We assessed the quality of distortion corrections produced by direct mapping to *Synth*-generated synthetic images in two primary contexts representing important use cases for fMRI data. First, we examined a subset of 100 participants selected randomly from the Adolescent Brain Cognitive Development (ABCD) study dataset (median age: 9.85 years; min: 9 years; max: 11 years). Our random sample used multi-band data acquired using GE, Philips, and Siemens scanners (see Table 1 in Supplemental Material). Second, we evaluated *Synth*'s performance on the Midnight Scan Club (MSC) dataset, which consists of resting state fMRI scans acquired from 10 participants on 10 separate occasions (300 min of resting state fMRI data/participant). The MSC precision functional mapping (PFM) (Braga and Buckner, 2017; Gordon et al., 2017, 2018, 2020; Gratton et al., 2018; Greene et al., 2020; Laumann et al., 2015; Lynch et al., 2020; Marek et al., 2018; Newbold et al.,

2020b,a; Sylvester et al., 2020; Zheng et al., 2020) dataset allowed us to assess session-to-session reliability of EPI distortion correction schemes across multiple sessions for the same participant. Importantly, image acquisition parameters for ABCD and MSC datasets differ significantly allowing us to assess the performance of *Synth* distortion correction on images with a range of contrast and levels of detail. Image acquisition parameters have been reported in detail elsewhere (Casey et al., 2018; Gordon et al., 2017).

In the ABCD and MSC datasets, we compared the effectiveness of this approach against five widely used distortion correction algorithms. Three of these approaches rely on separately acquired field map data (FSL *fugue*; FSL *topup*; and AFNI's 3dQwarp) while the fourth and fifth approaches implement alternative field map-less distortion correction methods, SyN-SDC (Esteban et al., 2019, 2021) and SynB0-DisCo. FSL field map corrections were constructed using either *fugue* (used for double echo field maps acquired in the MSC dataset) or *topup* (used for opposite direction phase encoded field maps in the ABCD dataset). In AFNI (Cox, 1996), the distortion correction is estimated from opposite phase encoded field maps using a "meet in the middle" nonlinear warp estimation implemented in 3dQwarp (Zhark the Grotesquely Warped (but still strangely handsome), 2021). This estimated warp is used to correct the image and combined with a separately estimated rigid body transform to align the functional to the anatomical image. Because 3dQwarp does not operate on double echo field map data, we could not assess its performance on the MSC data set. We therefore only compared *Synth*'s performance to FSL$_{fugue}$ and SyN SDC, SynB0-DisCo.

For field map-less approaches, we assessed the SyN-SDC field map-less method that corrects distortion by non-linearly aligning the participant's EPI and anatomical images while constraining allowable warps to regions known to be strongly affected by distortion (Huntenburg, 2014; Treiber et al., 2016; Wang et al., 2017). Additionally, we evaluated the SynB0-DisCo method, which uses deep learning methods to synthesize a pseudo infinite bandwidth EPI image as an input to FSL's *topup* to estimate the distortion correcting warp (Schilling et al., 2020).

Performance of the different distortion correction approaches was assessed in matched datasets, thus all statistical comparisons were performed within a multi-level modeling framework (fitlme, MATLAB). For these analyses, the metrics produced by the *Synth*-based registration pipeline were modeled as the baseline and differences in performance metrics associated with other approaches were modeled with main effect factors. To account for participant-specific variability, independent of the registration approach, participant identity was modeled as a random effect.

Studies in which the phase encoding direction is chosen to align with the anterior–posterior axis represent the most common use case. However, researchers have elected in some instances to choose other phase encoding directions for a variety of theoretical and practical reasons, for instance in an attempt to minimize total distortion or to preserve signal in orbitofrontal regions. The Human Connectome Project (HCP), in which the phase encoding direction was chosen to align with the left–right axis, represents a notable example of such a study (Van Essen et al., 2013; Glasser et al., 2016). While this approach is decidedly less common, it does represent an important use case. We therefore applied Synth to a 10 subject data set chosen at random from the publicly available HCP data set. We compared corrections produced by Synth against those produced by FSL's topup. Instances in which HCP results differ in character from those observed in the ABCD and MSC data set are noted in the relevant portion of this manuscript and are otherwise included in the supplement.

2.4. Evaluation metrics

No consensus exists as to what global metric best quantifies alignment quality. Many metrics used to measure alignment quality tend to bear some relation to the cost functions that are used during registration and therefore directly introduce a risk of confounding circularity. Thus, we computed several metrics comparing the similarity of different features of the field map corrected EPI images to their associated anatomical images.

2.4.1. Contrast similarity

To quantify the overall similarity between two aligned images from different modalities, we computed a contrast similarity metric which was defined as the linear correlation coefficient between vectorized versions of corresponding regions of the two images. Correlations were constrained to values that reside within a full brain binary mask.

2.4.2. Normalized mutual information (NMI)

A quantification of global image similarity determined by the amount of information shared between the time-averaged EPI image from each session and its associated T1w/T2w anatomical images (e.g., T1w-EPI NMI, T2w-EPI NMI).

2.4.3. Edge alignment

A quantification of alignment between any high-contrast edges existing in two images. It is defined as the correlation coefficient between the gradient magnitude images of the time-averaged EPI image and T1w/T2w images within a whole brain mask.

2.4.4. Segmentation alignment

In EPI images, gray matter, white matter, and CSF tend to exhibit distinct intensity values. When anatomical and EPI images are well aligned, segmentation maps derived from anatomical images and overlaid on the EPI image should correspond well to the tissue types of the EPI image. In well aligned images, the distributions of intensity values of EPI image voxels for a particular tissue type should tend to be more distinct and separable than for poorly aligned images. Here, separability is defined as the ability of a linear classifier (i.e. thresholding) to correctly distinguish between two tissue types. Each segmentation metric represents the area under the curve (AUC) of a receiver operating characteristic curve (ROC) generated from distributions of voxels delineated by two segmented tissue types. Each metric differs on the two tissue types selected. The segmentation of each tissue is limited to voxels 'adjacent' to the other tissue — where 'adjacent' refers to the set of all voxels in the 1st tissue that shares a face with a voxel in the 2nd tissue and vice-versa. AUC$_{gw}$ represents the separability of adjacent gray and white matter voxels, AUC$_{ie}$ represents the separability of adjacent voxels between the interior and exterior of the brain, and AUC$_{vw}$ represents the separability of adjacent voxels within ventricular cerebrospinal fluid and adjacent white matter.

2.4.5. Local alignment

Often, approaches to image registration rely on minimizing or maximizing a global distance or similarity metric (e.g., least-squares or mutual information) computed across an entire image volume. These metrics can be strongly influenced by the presence of intensity bias fields, which are low spatial frequency modulation of image intensity often caused by suboptimal participant placement or poor shimming. Additionally, experimental evidence has suggested that inaccurate registration can, in some situations, produce high values for many global measures of image alignment (Rohlfing, 2012). We therefore sought to compare the performance of each approach to EPI distortion correction with respect to a measure of local image similarity that would be less sensitive to these influences.

To do this, we implemented a 'spotlight' analysis examining 7x7x7 voxel regions (3 mm isotropic voxels) of the EPI images produced by each method and quantified its similarity to the corresponding region of the participant's T1w and T2w images with an R^2 metric. Image registration quality near gray matter is typically of greatest interest, hence we summarized local T1w-EPI and T2w-EPI similarity by computing the average spotlight R^2 across all gray matter voxels, as labeled by the participant's Freesurfer segmentation (Fischl, 2012).

2.4.6. Intra-participant stability

We used the MSC dataset to investigate how different distortion correction procedures affect metrics of inter-session stability for fMRI data. We reasoned that improving EPI image consistency would also reduce the variability of the intensities of individual voxels across sessions. To test this, we constructed time-average EPI images for each session contributed by a participant. We temporally concatenated each of these images, forming a pseudo-time series reflecting session-to-session changes in the EPI images. Then, we computed a metric, session-signal-to-noise (s-SNR, analogous to traditional t-SNR), for each voxel defined as the mean intensity of a voxel, μ, divided by the standard deviation of its intensity across sessions, σ. Lastly, we summarized alignment stability for each participant by averaging s-SNR across all voxels residing within a whole brain mask defined by their participant-specific FreeSurfer parcellation.

Operating on the assumption that less reliable field map correction procedures would introduce additional session-to-session variability in resting state functional connectivity (RSFC) matrices (pairwise correlations between voxel time series), we measured the similarity of each pair of RSFC matrices contributed by a given participant. We identically preprocessed (see Appendix A.2.4 in Supplemental Methods) the MSC datasets produced by each registration pipeline and computed the RSFC BOLD-signal correlation matrices for all gray matter voxels. For each session, we extracted the upper triangle of the associated correlation matrix. Finally, we quantified resting state correlation matrix stability by computing the pairwise correlation between the upper triangles of each unique pair of a participant's resting state matrices.

3. Results

3.1. Synth images consistently match the contrast of EPI images

We reasoned that anatomically-based reference images with greater contrast similarity to EPI images would serve as a more reliable reference for distortion correction than images with lower contrast similarity. To assess differences in contrast similarity between EPI and T1w, T2w and synthetic images, we calculated the linear correlation coefficient between participants' EPI images and each of their T1w, T2w and synthetic images within a whole brain mask. By default, *Synth* uses a model that includes a 12-component radial basis function (RBF) decomposition of the T1w/T2w source images along with pair-wise T1w/T2w interaction terms to create synthetic EPI images (Fig. 1). This model was sufficient to produce synthetic images with comparable EPI contrast similarity in both ABCD and MSC datasets (Fig. 2a,b). In both datasets, synthetic images exhibited significantly increased contrast similarity to EPI images compared to T1w and T2w images (Fig. 2c). Contrast similarity between the EPI and *Synth* images was consistently highest. Mixed-effects model comparisons in the ABCD data set revealed that the differences between *Synth* contrast similarity and T1w/T2w contrast similarity were both significant: (T1w-*Synth*= −0.417; $p < 0.001$; $t = -46.18$; $df = 297$) and (T2w-*Synth*= −0.18; $p < 0.001$; $t = -19.62$; $df = 297$). The same pattern held in the MSC dataset: (T1w-*Synth*= −0.374; $p < 0.001$; $t = -49.64$; $df = 297$) and (T2w-*Synth*= −0.274; $p < 0.001$; $t = -36.4$; $df = 297$). Synthetic EPI images derived from the HCP data set exhibited greater contrast similarity to their respective EPI reference images than their associated T1w/T2w images (Supplemental Figure 4)

3.2. Synth outperforms existing field map-less methods on global measures of image alignment

We assessed *Synth*'s distortion correction performance (Fig. 3) against existing field map-based (AFNI 3dQWarp; FSL *topup*; FSL *fugue*) and field map-less (SyN-SDC; SynB0-DisCo) distortion correction methods using both established and novel metrics that quantify global image similarity between the EPI image and the associated T1w/T2w

anatomical images (Fig. 4). These metrics included normalized mutual information; the correlation coefficient between the image gradients; and segmentation alignment. Here, we outline *Synth* registration performance for each metric described above as compared to the highest performing competitor. *Synth* registration tended to produce EPI images with the highest registration quality metrics for both ABCD and MSC data. Comparisons are reported in the form of statistical contrasts corresponding to the mean difference between registration metrics (e.g., Method A-Method B). Full statistical tables comparing all methods are included in section 5.4 and 5.5 of the Supplemental Materials. Illustrative comparisons between FSL fugue/topup and Synth are included in Supplemental Figures 1 and 2.

First, we examined a well-established image similarity metric, normalized mutual information (NMI). In the ABCD dataset, NMI between EPI and T1w images was not significantly different (FSL_{topup}-*Synth*= 9.25e-6; $p = 0.98$; $t = 0.025$; $df = 495$). In contrast, NMI for T2w images was significantly higher for *Synth* images than those aligned by FSL_{topup} (FSL_{topup}-*Synth*= −1.49e-3; $p < 0.001$; $t = -2.85$; $df = 495$). *Synth* alignment produced EPI images with the highest NMI shared between T1w/T2w images in the MSC dataset as well (FSL_{fugue}-*Synth*= −4.02e-3; $p < 0.001$; $t = -10.74$; $df = 396$) and (FSL_{fugue}-*Synth*= −3.22e-3; $p < 0.001$; $t = -8.53$; $df = 396$) (Fig. 4a). We observed no significant differences in these alignment metrics in the HCP data set (Supplemental Figure 5, Supplemental Section 5.6).

Next, we examined the quality of alignments emphasizing registration of high contrast boundaries. EPI-T1w edge alignments produced by *Synth* were outperformed by those produced by SyN SDC (SyN SDC-*Synth*= 3.19e-2; $p < 0.001$; $t = 4.93$; $df = 495$) in the ABCD dataset; for T2w gradient magnitude images, there was no significant difference between *Synth* and the highest performing competitor, FSL_{topup} (FSL_{topup}-*Synth*= −3.12e-3; $p < 0.001$; $t = -0.54$; $df = 495$). In the MSC dataset, its closest competitor varied by modality. For EPI-T1w edge alignment, FSL_{fugue} was the highest performing competitor, but still produced edge alignments that were significantly lower than those produced by *Synth* distortion correction (FSL_{fugue}-*Synth*= −1.70e-2; $p < 0.001$; $t = -3.22$; $df = 396$). For T2w images in the MSC dataset, FSL_{fugue} was once again the closest competitor but did not perform significantly better than *Synth*, (FSL_{fugue}-*Synth*= 4.88e-3; $p = 0.39$; $t = 0.87$; $df = 396$) (Fig. 4b). In the HCP data set, these metrics revealed no significant differences between *Synth* and FSL_{topup}.

In addition, we examined the quality of the alignment of EPI images to anatomical images segmented by tissue type (gw: gray/white matter; vw: ventricles/white matter; ie: interior/exterior of brain) based on the participants' freesurfer segmentation. These metrics quantify the separability of EPI image voxel intensities based on anatomically defined segmentation (higher AUC [area under the receiver operating characteristic curve] is better). AUC_{gw} image segmentation alignment metrics for *Synth* were significantly lower than those produced by FSL_{topup} in the ABCD dataset (FSL_{topup}-*Synth*= 1.62e-2; $p < 0.001$; $t = 5.86$; $df = 495$) and MSC (FSL_{topup}-*Synth*= 1.31e-2; $p < 0.001$; $t = 6.03$; $df = 396$) datasets. *Synth*-aligned images exhibited greater AUC_{ie} in both ABCD (FSL_{topup}-*Synth*= −1.60e-2; $p < 0.001$; $t = -5.60$; $df = 495$) and MSC (FSL_{fugue}-*Synth*= −2.16e-2; $p < 0.001$; $t = -8.82$; $df = 396$) datasets. We also found that the SyN SDC pipeline produced the greatest AUC_{vw} values (SyN SDC-*Synth*= 2.14e-2; $p < 0.001$; $t = 3.77$; $df = 495$) for the ABCD dataset. However, *Synth* performed better than the highest performing competitor, FSL_{fugue} (FSL_{fugue}-*Synth*= −1.02e-2; $p = 0.03$; $t = -2.21$; $df = 396$) in the MSC dataset (Fig. 4c). In the HCP data set, FSL_{topup} produced significantly greater AUC_{gw} values (FSL_{topup}-*Synth*=1.63e-2; $p < 0.001$; $t = -4.52$; $df = 18$) whereas corrections produced by *Synth* exhibited greater AUC_{vw} values (FSL_{topup}-*Synth*=−1.29e-2; $p = 0.001$; $t = -2.82$; $df = 18$)
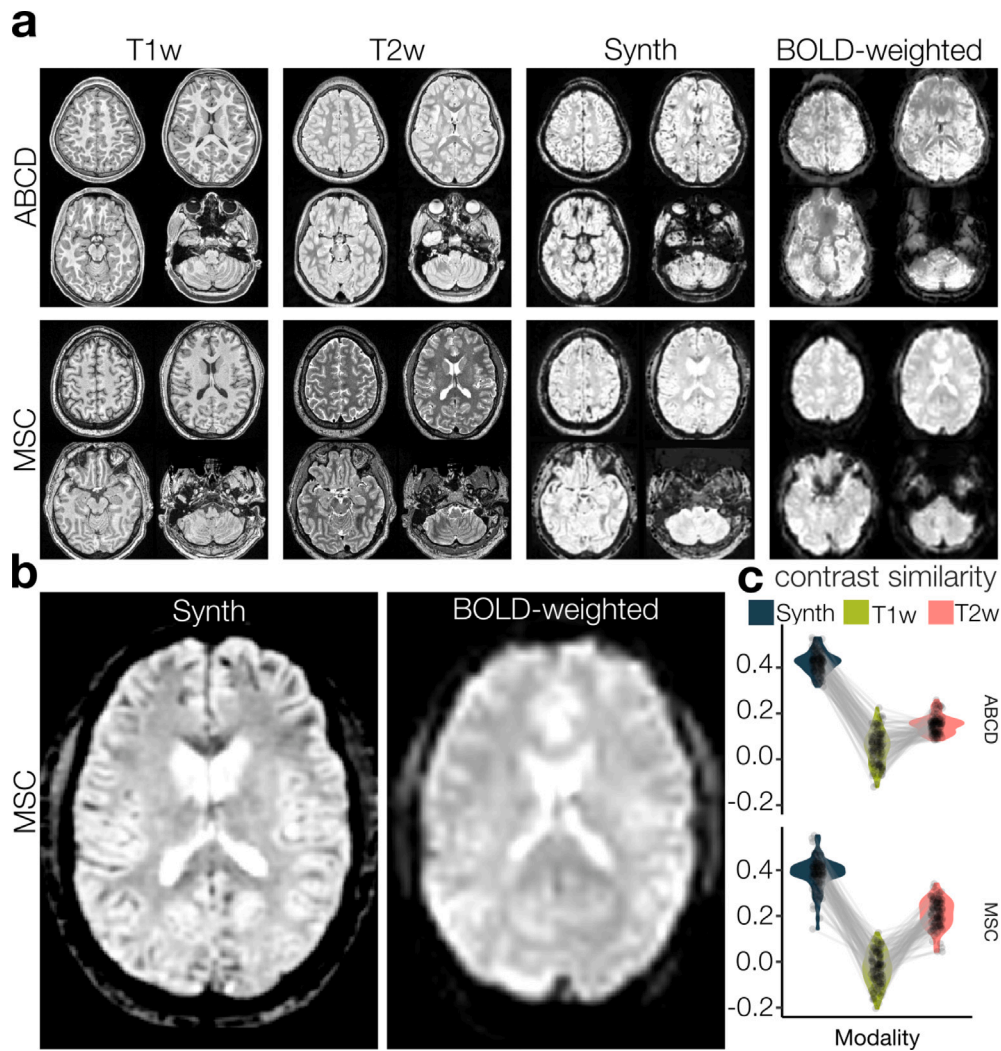
Fig. 2. Image contrast similarity between T1w, T2w, *Synth* and EPI images. (a) T1w, T2w, synthetic, and corresponding affine aligned EPI images (sample participants; ABCD top, MSC bottom). **(b)** Enlarged view of an axial slice of the Synthetic (left) and real EPI (right) images (MSC example). **(c)** Average contrast similarity between EPI images and associated anatomical and *Synth* images across 100 participants from the ABCD dataset (upper). Average contrast similarity between a true EPI image and associated anatomical and synthetic functional images for each participant in the MSC dataset (lower; 10 scans per participant).
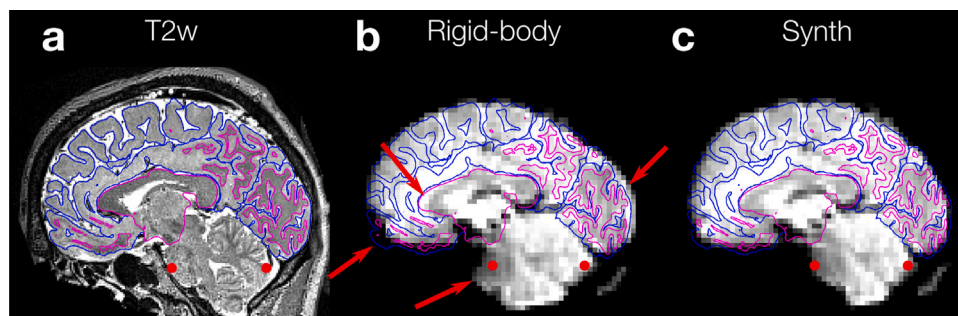
Fig. 3. Example of distortion correction with *Synth* compared to rigid-body alignment (a) Parasagittal slice of T2w image from an example MSC participant. Fuschia and blue lines indicate gray/white matter boundary estimated by freesurfer segmentation of the associated T1w image. **(b)** Corresponding slice of functional EPI image aligned to the anatomical image using rigid body alignment procedure implemented in FSL's FLIRT. Red arrows indicate regions of local misalignment due to image distortion. Red dots indicate anatomical fiducial markers to aid in assessing distortion correction applied to cerebellum and pons. **(c)** Identical slice through EPI image corrected using *Synth*. Note improved local alignment of corpus callosum, occipital pole, pons, and orbitofrontal cortex. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
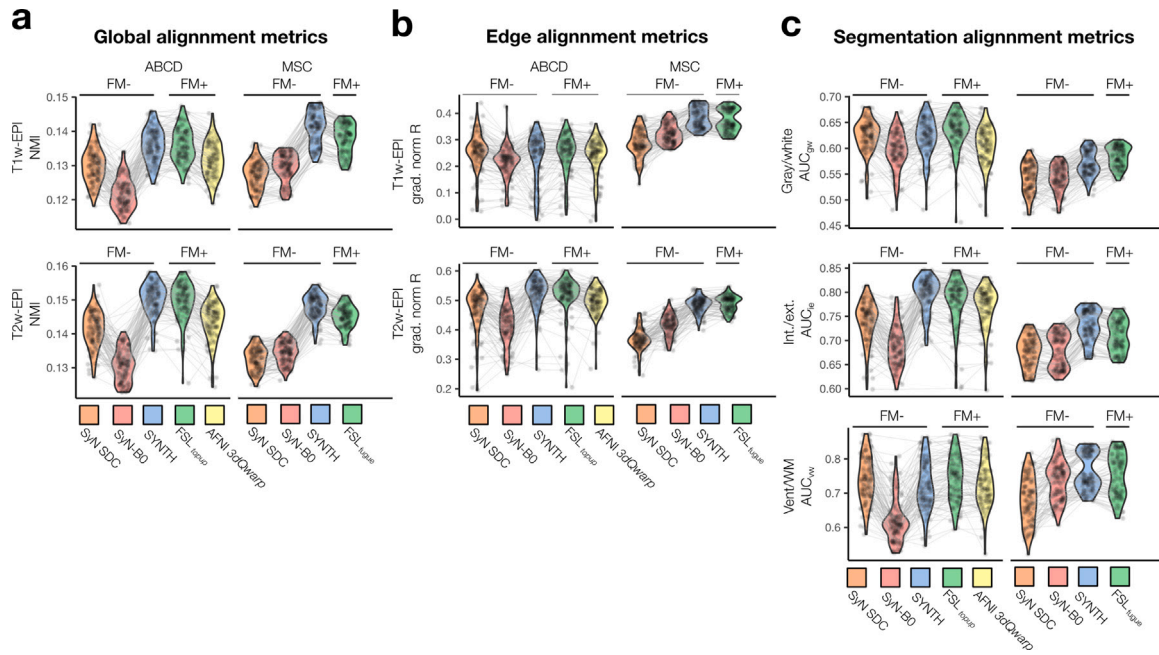
Fig. 4. Alignment metric comparisons for distortion correction pipelines. Violin plots depict the distributions of each alignment metric for ABCD (left column) and MSC (right column) datasets. Horizontal bars group assessment metrics for field map-less distortion correction methods (FM-) and field map based (FM+) distortion correction methods. **(a)** Global similarity: Normalized mutual information (NMI) shared between the registered EPI images and their associated T1w/T2w anatomical images to assess global similarity. **(b)** Edge alignment: Correlation coefficient between the gradient magnitude images for EPI and T1w/T2w images for a metric of edge alignment. **(c)** Segmentation alignment: Metrics that quantify the separability of anatomical features in EPI data base on the Freesurfer segmentation of their anatomical T1w images.

3.3. Synth improves local measures of anatomical and functional image alignment

Existing experimental evidence indicates that field map correction quality varies regionally and depends on the method used to estimate the B0 inhomogeneity (Schallmo et al., 2021). Consequently, to identify whether *Synth* was better at correcting certain brain regions, we carried out a winner-take-all (WTA) analysis across all voxels by determining which method produced the greatest average local similarity between T1w and T2w anatomical images (R^2 computed within a 7x7x7 voxel spotlight). Representative spotlight R^2 images underlying the WTA analysis are provided in Supplemental Figures 6 and 7. Each distortion correction method exhibited regions of the brain for which they performed consistently better than other distortion correction methods. Although the structure of the WTA maps varied between datasets, they shared some consistent features (Fig. 5).

In ABCD and MSC datasets, *Synth* produced either the largest, or second largest regions of high quality local alignment, regardless of anatomical image modality (i.e., T1w, T2w). In the MSC data set, *Synth* generally had the best regional correction performance across the entire brain. For the ABCD dataset, the best regional distortion correction performance was split between FSL$_{topup}$ and *Synth*, where FSL$_{topup}$ provided better local alignment to T1w images over a large swathe of the superior frontal cortex while *Synth*'s particular strengths were in improving local registration of cerebellar, brainstem, orbitofrontal and occipital regions with respect to a participant's T2w image. Overall, FSL$_{topup}$ won the greatest percentage of voxels in the T1w and T2w comparisons for the ABCD dataset, (T1w: FSL$_{topup}$= 43%; *Synth*= 33%; T2w: FSL$_{topup}$= 51%; *Synth*= 46%). While *Synth* won the greatest percentage of voxels in the T1w and T2w comparisons over FSL$_{fugue}$ (T1w: FSL$_{fugue}$= 21%; *Synth*= 65%; T2w: FSL$_{fugue}$= 31%; *Synth*= 54%) for the MSC dataset. In the HCP data set, FSL$_{topup}$ produced better measures of local alignment over a modestly larger fraction of the image (T1w: FSL$_{topup}$ = 55%; *Synth*=45%; T2w: FSL$_{topup}$ = 57%; *Synth*=43%) (See Supplemental Figures 8, 9, and 10). These results indicate that *Synth* can perform at parity with state-of-the-art field map-based approaches.

The analysis of functional images depends to a greater extent on the quality of gray matter alignment and less so on the alignment quality of other voxels representing other tissue types. Accordingly, we summarized local gray matter alignment performance of each distortion correction method by computing the average spotlight R^2 across only the gray matter voxels. We observed that *Synth* and FSL$_{topup}$ registration pipelines produced similar quality of local registration to the participants' T1w image (FSL$_{topup}$-*Synth*= -2.51e-3; $p = 0.17$; $t = -1.38$; $df = 495$), in the ABCD dataset. *Synth* performance did not differ significantly from the highest performing competitor, FSL$_{topup}$, for average local similarity between EPI and T2w images (FSL$_{topup}$-*Synth*= -4.14e-3; $p = 0.21$; $t = -1.27$; $df = 495$). *Synth* registration produced EPI images with the greatest average local similarity for associated T1w images (FSL$_{fugue}$-*Synth*= -1.93e-2; $p < 0.001$; $t = -8.64$; $df = 396$) and did not differ significantly from the highest performing competitor, FSL$_{fugue}$ for T2w images (FSL$_{fugue}$-*Synth*= -3.27e-3; $p = 0.17$; $t = -1.37$; $df = 396$) in the MSC dataset (Fig. 6).

3.4. Synth improves intra-participant consistency across sessions

The primary goal of applying distortion correction to EPI images is to improve data quality by reducing variability caused by poor alignment. It is therefore important to know the extent to which a particular approach to distortion correction introduces additional measurement variability. By applying different distortion correction strategies to the repeated-measures in the MSC dataset, we assessed how each method contributed to within-participant structural and functional variability.

First, we examined each voxel's average intensity across all sessions normalized by its standard deviation across sessions. We term this stability metric, session-to-session SNR (s-SNR). We observed that for *Synth*-aligned images, s-SNR was not significantly different than that produced by the highest performing competitor, FSL$_{fugue}$ (FSL$_{fugue}$-*Synth*= 0.80; $p = 0.55$; $t = 0.60$; $df = 36$) (Fig. 7a).

Next, we quantified global similarity of time averaged EPI images for each pair of sessions contributed by a participant. We extracted the region of the aligned time-averaged EPI volumes containing just brain
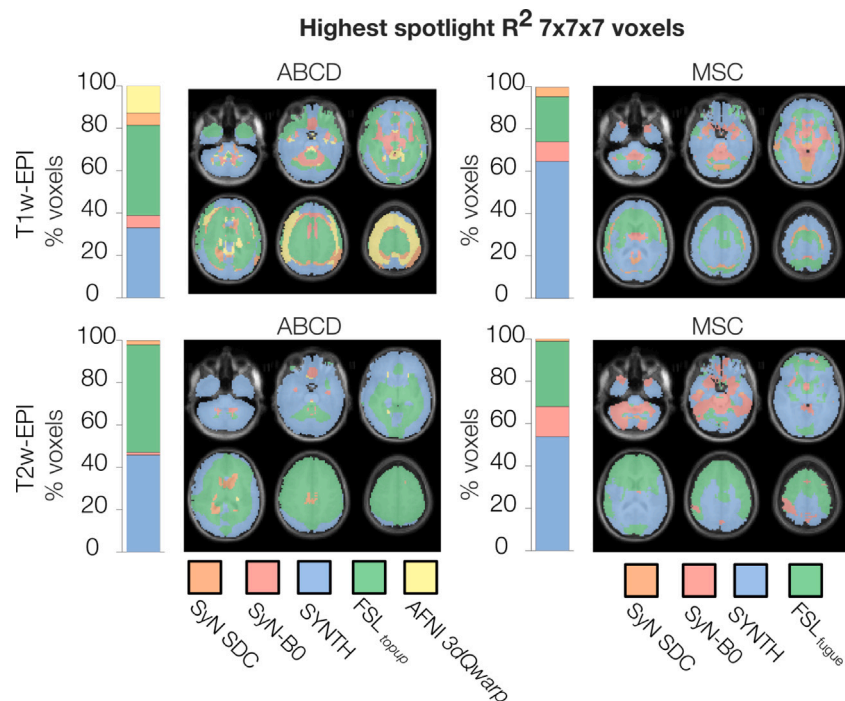
Fig. 5. Regional variation in distortion correction performance. Summary of local similarity between the EPI images produced by each registration pipeline and their associated T1w and T2w anatomical images. Winner-take-all maps depicting, for each voxel, the pipeline that produced the greatest average local correlation between corresponding regions of EPI and T1w/T2w images. The color of each voxel depicts which registration pipeline produced, on average, the greatest spatial R2 values computed over a centered 7x7x7 cube. Bars to the left of each set of images depict the percentage of voxels "won" by each registration pipeline. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

matter (as classified by the FreeSurfer segmentation of the participant's anatomical images) from each session. Then, for each participant, we computed the average linear correlation coefficient across all pairs of sessions. Owing to the general stability of a given participant's brain geometry on the time scales over which the MSC dataset was collected, inter-session linear correlation coefficients between scan sessions tended to be quite high across all methods ($0.94 - 0.99$). While *Synth* produced EPI images with the greatest session-to-session similarity, it did not produce images that were significantly more stable than those produced by FSL_{fugue} (FSL_{fugue}-*Synth*$= -5.56\text{e-}3; p = 0.23; t = -1.20; df = 1796$) (Fig. 7b).

Finally, we examined how the choice of distortion correction approach affects the stability of resting state functional connectivity (RSFC) data. To do this, we created voxel-level RSFC matrices for each session of data and computed RSFC similarity between pairs of sessions as the correlation between upper triangular portions of the RSFC matrices. Mixed effects analysis of similarity between unique session pairs revealed RSFC correlation matrices from *Synth* aligned datasets were not significantly less stable than those produced by the highest performing competitor, FSL_{fugue} (FSL_{fugue}-*Synth*$= 5.47\text{e-}3; p = 0.12; t = 1.55; df = 1796$) (Fig. 7c).

4. Discussion

Distorted fMRI images are an inevitable consequence of the trade-offs made to achieve a sampling rate sufficient to measure neurally meaningful fMRI signal variability in Cartesian sampled EPI (Glover, 2012). Complications in registering distorted functional EPI images to their undistorted anatomical image counterparts are a perennial challenge faced by neuroimaging researchers. Performed well, proper image registration reduces variability in measurements, improves the statistical power of analyses, and produces significant improvements in localization of responses to experimental manipulations (Devlin et al., 2000; Togo et al., 2017; Visser et al., 2010a,b). For these reasons, the quality of EPI distortion correction and anatomical registration

is of great concern in every branch of neuroimaging and the primary motivation for developing methods to improve fMRI registration performance.

We hypothesized that the reliability of field map-less approaches could be improved if the underlying contrast properties of the EPI and anatomical images match more closely. The distortion correction approach that *Synth* employs is to generate a high resolution synthetic functional image that exhibits greater similarity to tissue contrast properties of a EPI image but is based on the information from the undistorted geometry of a participants' T1w and T2w anatomical images. The approach implemented in *Synth* for constructing a synthetic EPI image target can be considered an example of a class of registration procedures termed "mono-modal reduction" (Sotiras et al., 2013). After constructing a synthetic target image, a distortion-correcting warp is constructed in the same way as other direct mapping approaches: by directly mapping to an undistorted target image. Here, we have demonstrated that when a synthetic image is used as a target, the resulting fMRI alignment quality rivals and in some cases exceeds that produced by standard field map methods.

4.1. Synth achieves parity with or exceeds the distortion correction performance of alternative methods on most measures of global alignment

We began our exploration of distortion correction performance by examining the commonly used global image similarity metric, normalized mutual information. We found that *Synth* performed comparably to or better than state of the art field map corrections provided by FSL's *topup* and *fugue* in both ABCD and MSC data sets. One of the challenges to assessing the quality of multi-modal image registration and distortion correction is determining which global metric serves as the optimal summary of final alignment quality. To address this concern, we assessed the performance of each distortion correction approach with several additional global image similarity metrics, each intended to assess alignment with respect to different image features.
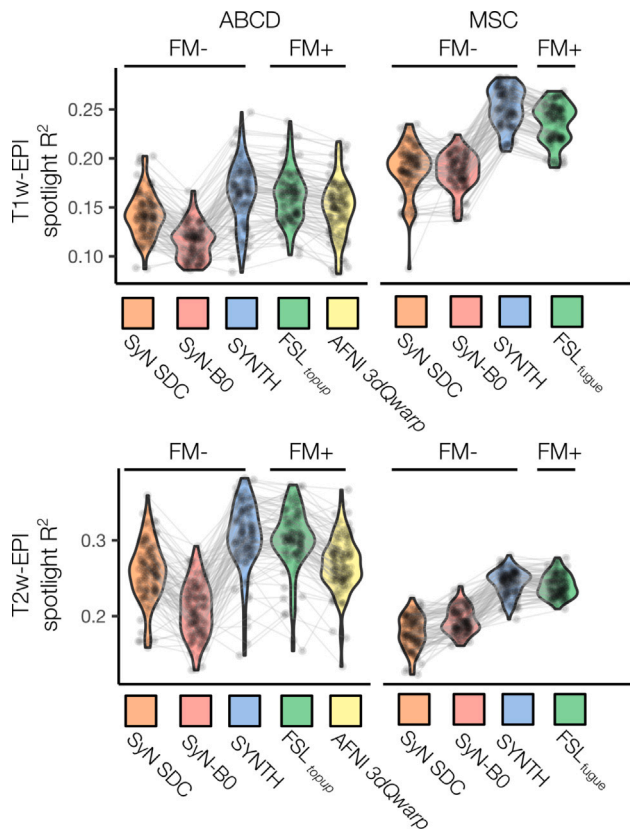
Fig. 6. Average local image similarity across all gray matter voxels Summary of local similarity between the EPI images produced by each registration pipeline and their associated T1w/T2w anatomical images. Each plot depicts the distributions of average spotlight R2 value between each participant's EPI image and T1w/T2w images computed across all gray matter voxels.

The first alternative global similarity metric we examined was the quality of high-contrast edge alignment. We reasoned that these features would be less influenced by residual low spatial frequency bias fields and the differences in tissue contrast properties between images acquired with different modalities (T1w, T2w, EPI). Here, *Synth*'s performance on cross-modal edge alignment metrics was generally high, either achieving the highest performance or placing in the top two. The single exception to this outcome was observed in the T1w-EPI edge alignment metrics in the ABCD data set. There, FSL's *topup* and SyN SDC produced the highest alignment metrics. This result may arise from the fact that the *Synth* pipeline initially registers the EPI image into the anatomical image space using a T2w reference image, while the FSL and SyN SDC pipelines both register to the anatomical image space using the T1w image as a reference. It may be the case that the choice of image modality (i.e. T1w/T2w) used for the initial EPI-anatomical alignment biases the edge alignment metric for those two image modalities. Significantly, we observed that *Synth*'s high performance generalized well to the HCP data set in which, due to the left–right orientation of the phase encoding direction, EPI distortion exhibits very different structure.

A second notable exception to *Synth*'s typically high performance was the outcome of the segmentation alignment metric focusing on the alignment accuracy of the ventricles and white matter (AUC_{vw}) in the ABCD dataset. In that instance, the field map-less SyN SDC method provided the best ventricles and white matter alignment metrics (AUC_{vw}) in the ABCD dataset followed by FSL's *topup*. One potential reason for this outcome is that both FSL's *topup* and SyN SDC pipelines used brain-based registration for their EPI to anatomical alignments (Greve and Fischl, 2009). This registration technique incorporates anatomical

segmentation information when guiding anatomical alignments and may improve ventricle/white matter correspondence between the EPI and anatomical images. This interpretation is complicated somewhat by the fact that ventricular alignment metrics were greatest for *Synth*-based alignments in the HCP data set, while FSL's *topup* tended to produce better gray/white matter alignment.

4.2. Synth exhibits high performance in regional measures of image alignment

In addition to the challenge of selecting and justifying a particular global metric for measuring multi-modal image alignment quality, prior work has shown it does not always provide a sufficient condition for demonstrating optimal alignment (Greve and Fischl, 2009). This is because global metrics, while flexible, often cannot fully capture the complicated mapping of voxel intensity values between images acquired with different modalities. While achieving high measures of global image similarity is a necessary condition for quality image alignment, it is important to consider them alongside non-global measures to fully assess alignment (Saad et al., 2009). For this reason, we examined small regions of the image (i.e. spotlight), which contain a limited range of tissue types and minimal residual bias field. In this way, we created an alignment metric that assesses image similarity while reducing the influence of non-linearities in the relationship between voxel intensity values of images acquired with different modalities.

Our winner-take-all analyses of local image alignment revealed *Synth* to be among the top performing distortion correction approaches, producing the highest local image quality metrics in large swathes of the cortex, but also performing particularly well in correcting distortions affecting the cerebellum and brain stem. That different distortion correction approaches excelled in correcting different areas of the brain is in line with existing evidence suggesting that the accuracy of standard field map distortion correction methods varies across the brain (Schallmo et al., 2021; Graham et al., 2017).

One possible reason for *Synth*'s superior performance in the MSC dataset overall, compared to the ABCD dataset, may be due to the dual echo field maps used in the MSC dataset. Prior work has shown that distortion correction with reverse-phase encoding field maps outperform corrections produced by dual echo field maps (Graham et al., 2017). Our findings are consistent with these observations, and suggest that *Synth* may provide superior correction to dual-echo field map corrections in datasets in which image acquisition parameters share sufficient similarity to those used for the MSC study (Gordon et al., 2017).

4.3. Synth does not introduce session-to-session variability like existing field map-less approaches

We observed that as a general rule, field map-less approaches tended to introduce more session-to-session structural variability than traditional field map-based approaches. *Synth* stood out as a clear exception to this trend by introducing no more session-to-session variability than a simple rigid-body alignment procedure or field map correction. Importantly, because *Synth* distortion correction did not measurably increase the session-to-session variability of RSFC matrices we can conclude that the improvement in EPI-to-anatomical alignment provided by *Synth* does not come at a hidden cost of increasing the variability of the underlying functional data.

4.4. Accounting for differing performance of existing field map-less approaches

Synth tended to produce higher quality and more reliable corrections than existing field map-less approaches (SyN SDC and SynB0-DisCo) with which it shares many conceptual similarities. This prompts
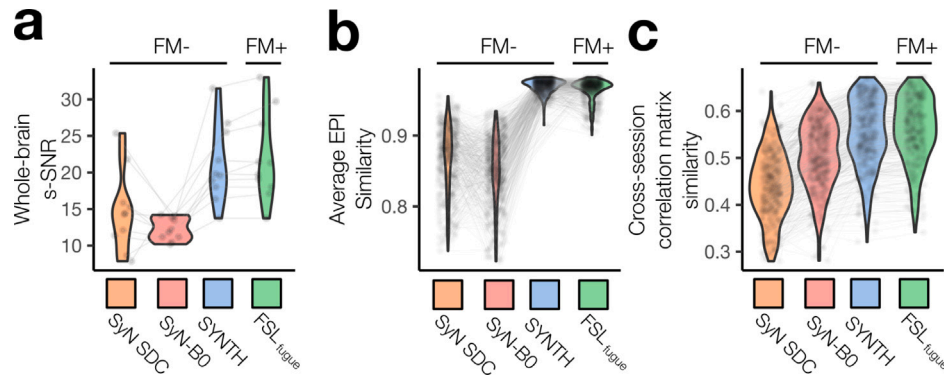
Fig. 7. Cross-session structural and functional reliability metrics. Repeated measurements of the same 10 individuals in the MSC dataset allows us to assess the reliability of different approaches to EPI distortion correction. **(a)** Average voxelwise session-SNR (s-SNR) computed within a whole brain mask. s-SNR was computed as $\frac{\mu}{\sigma}$, where μ is the voxel intensity of the time average EPI image averaged across all sessions, and σ is the standard deviation. **(b)** Distributions of the mean pairwise correlation between time-averaged resting state EPI volumes. **(c)** Similarity between the EPI signal correlation matrices from each unique pair of sessions. Each of the 10 participants contributed 10 resting state datasets, corresponding to 45 unique session pairs per participant. Similarity is quantified as the correlation between the vectorized upper triangles of the pairs of EPI signal correlation matrices.

us to consider the potential reasons why *Synth* performs with greater reliability and accuracy.

In the case of SyN SDC, one reason for decreased reliability may be the use of a group average field map template in order to constrain its solutions. Such a template may not allow sufficient flexibility to align high spatial resolution participant-specific features. SyN SDC also relies on an intermediate atlas alignment stage, which may be a source of additional variability. Because participant motion can influence the structure of the EPI distortion, the use of a field map template that does not account for this effect may reduce overall efficacy.

Both *Synth* and SyN SDC employ a direct mapping approach to correct EPI distortion. The fundamental difference in their approaches is that SyN SDC attempts to directly align to a participants' T1w image –whose contrast properties are very different from the EPI image– while *Synth* uses an intermediate synthetic image whose contrast properties are more similar to the EPI. That *Synth* proved more effective and reliable than the similar approach implemented in SyN SDC would seem to indicate the importance of matching the contrast properties of target and reference images when attempting to use direct mapping approaches for distortion correction.

With SynB0-DisCo, a synthetic pseudo-infinite bandwidth intermediate EPI image facilitates distortion correction using FSL's *topup*. The corrections, though, were less accurate and less reliable than those produced by *Synth* and SyN SDC. On the surface, it would seem that the use of a contrast matched intermediate image along with FSL's state of the art field map estimation software should produce the same reliable distortion correction observed with *Synth*. Two factors may potentially contribute to the variable performance of this approach. The tissue contrast properties of EPI data can vary significantly depending on image acquisition parameters. Therefore, one possibility is that SynB0-DisCo's deep learning model may generalize poorly to unseen datasets whose resolution and contrast are quite different from its training data. A second possibility is that synthetic images produced by deep learning models are prone to introducing spurious artifacts into their outputs (Antun et al., 2020; Bhadra et al., 2021). The presence of such artifacts may be a significant source of registration variability. Both of these effects may have contributed to reducing SynB0-DisCo's performance in the ABCD and MSC datasets. *Synth*'s ability to generate effective synthetic images based solely on the participant's anatomical data without relying on a large training dataset avoids both of these potential issues by flexibly adapting to the contrast properties of a specific acquisition and minimizing the possibility of spurious artifacts.

While our results indicate that *Synth* is able to correct distortion comparably well to current field map based techniques, it is possible that advances in theory or software may improve field map-based corrections even further so that they reliably exceed *Synth*'s performance.

We therefore do not advocate abandoning the collection of field map data entirely. Rather we propose that *Synth* distortion correction is a reliable substitute for field maps at the present time and given the current state of the art. We can unreservedly recommend using *Synth* as an effective tool for correcting distorted fMRI data when no field map data is available.

4.5. General guidelines for deploying Synth

Our results indicate that synthetic images with contrast similarity (linear correlation coefficient) of 0.4 with the target EPI image are sufficient for producing high quality corrections. In order to accommodate datasets with different contrast, resolution, and fields-of-view, *Synth* provides a variety of command line options: *source_images*, is a list of pre-aligned undistorted anatomical images from which the target image is modeled. Typically this is the participant's T1w and T2w images, but any undistorted images can be used in principle; *interaction_model* is a string that specifies both main-effect terms and interaction terms between a user-specified number of source images, each decomposable into a user-specified number of RBF components. The number of RBF components applied to each source image enables more complex models that will allow *Synth* to capture more complicated relationships between the voxel intensities of the source images and the EPI images. This flexibility comes at the cost of increased memory requirements and computational time, but can produce more accurate synthetic images that may ultimately improve registration quality. In our results, the RBF model parameters were chosen to produce a synthetic image of qualitatively sufficient visual similarity to their associated EPI images. By default, *Synth* uses the RBF model parameters optimized for the ABCD and MSC datasets (i.e., a 12 component decomposition on each of the participants' T1w/T2w images, along with pairwise interaction terms between T1w/T2w components). In general, *Synth* attempts to map the contrast of high resolution source images to a low resolution EPI image, which requires modeling the difference in point spread function between the two. Modifying the bandwidth parameter enables *Synth* to be applied to datasets acquired at different resolutions. Users acquiring high resolution EPI images will decrease this parameter and users acquiring low resolution EPI images will increase it. Lastly, the *weight_mask* parameter is a binary mask that indicates the region of the EPI image in which a valid mapping between source image (T1w/T2w) intensities and target image (EPI) intensity values exists. This mask will exclude areas of extreme signal dropout. Only voxels within this mask are included when estimating the model that produces the synthetic image.

We provide a reference pipeline implementing *Synth* which only requires the user to organize their data into the BIDS format standard (Gorgolewski et al., 2016). Using this pipeline, all of *Synth*'s parameters will be automatically determined. Users seeking to integrate *Synth* into their own pipelines, will need to construct their own weight masks and determine the appropriate Epanachnikov smoothing kernel bandwidth independently. In general, we recommend that inputs to *Synth* should be deobliqued and bias field corrected, but otherwise minimally preprocessed. For walkthroughs detailing the *Synth*'s installation and use as a standalone utility or as part of our BIDS compliant reference pipeline, users are invited to explore https://gitlab.com/vanandrew/omni.

4.6. Future directions for field map-less distortion correction

Although the distortion corrections produced by *Synth* are high quality, several open questions remain relating to how this general approach can be implemented most effectively. Chief among these is whether *Synth*'s performance can be augmented further by including existing field map information derived using traditional approaches. The focus of the present report is assessing the quality of the field map corrections produced by *Synth* in a situation in which there is no initial estimate of the EPI image distortion provided to the underlying nonlinear warping software. Including an existing field map as an initial estimate may further increase *Synth*'s performance. Future iterations of *Synth* will include this ability and may produce even greater correction fidelity between functional and anatomical images than presented here.

Even though *Synth*'s underlying RBF model provides a great degree of flexibility in fitting the functions that map T1w/T2w voxel intensities to EPI image voxel intensities, other data collection parameters and preprocessing steps may improve the quality of the synthetic images or reduce the needed complexity of the RBF model. For example, adjusting echo times during data acquisition may produce images that can be more accurately modeled with fewer RBF components while still retaining sufficient sensitivity to BOLD contrast. Alternatively or in addition, performing contrast enhancing preprocessing of the anatomical and EPI images so that the resulting function relating their voxel intensities can be modeled more efficiently by the chosen RBF model may allow for simpler RBF models to perform equivalently. In our own preliminary tests and for the results reported here, we observed that synthetic images were more accurate when a bias field was estimated and removed from the anatomical images and image contrast was increased through histogram normalization (see Appendix A.2.2 in Supplemental Methods).

4.7. Conclusion

The results reported here have demonstrated that it is possible to achieve high-quality EPI distortion correction without the need to collect separate field map data. We have shown that field map-less approaches, such as *Synth*, can perform comparably to existing gold-standard distortion correction approaches, and in some cases, may surpass them on measures of global and local image alignment quality. Importantly, *Synth* produced high quality alignments even within the HCP dataset, indicating that its high performance generalizes not only across differences in EPI resolution and contrast properties, but is robust to differences in phase encoding direction. Removing the reliance on field maps to correct EPI distortion will allow researchers to recover samples with missing or corrupted field maps while maintaining high quality alignment between their anatomical and EPI images. Field map-less distortion correction may prove to be a particular asset to researchers studying high motion cohorts, such as pediatric or neuropsychiatric populations (Greene et al., 2016), where acquiring high quality field map data is a significant challenge. Reliable field map-less distortion correction shows great promise for overcoming the limitations arising from the acquisition, processing, and quality checking of field map data and has the potential to greatly simplify data processing in the neuroimaging field.

Code sharing

The *Synth* software package and other utilities and scripts used for this project can be downloaded at https://gitlab.com/vanandrew/omni.

Declaration of competing interest

Data availability

The ABCD data repository grows and changes over time. The ABCD data used in this report came from ABCD collection 3165 and the Annual Release 2.0, http://dx.doi.org/10.15154/1503209.

The MSC dataset used in this report can downloaded at can be found on OpenNeuro.org, http://dx.doi.org/10.18112/openneuro.ds000224.v1.0.3.

HCP data can be downloaded at https://humanconnectome.org/study/hcp-young-adult.

Acknowledgments

Appendix A. Supplemental methods

A.1. Description of common pipeline

Distortion correction methods were compared using a common registration pipeline, which differed only in the distortion correction method and functional-to-T1w alignment. This includes alignment of anatomical images to atlas, functional framewise alignment and re-sampling of the final aligned functional data to 3 mm. Details of each common step are explained in the following sections.

A.2. HCP dataset

Ten randomly selected participants from the Human Connectome Project data set were used for the supplemental analyses detailed in this manuscript. EPI data was drawn from the "Emotional Processing" portion of the dataset, along with each participants' T1w/T2w images. A thorough account of the acquisition details for task and anatomical imaging data are located in Barch et al. (2013) and Van Essen et al. (2013). Briefly, data was acquired using a modified SIEMENS Skyra using a TR of 720 ms, TE 3.31 ms, flip angle of 52 degrees and multiband acceleration factor of 8. EPI images were acquired with 2.0 mm isotropic voxels.

A.2.1. Deobliquing

Prior to any preprocessing, all data is deobliqued by modifying the orientation information of the NIFTI headers. This is accomplished by decomposing each orientation affine into its component translation, scale, rotation, and shear parameters and recomposing the orientation affine without the rotation component (Brett, 2023). The data is re-saved using the NiBabel library, which correctly sets the s-form/q-form of the NIFTI header for reading by other neuroimaging tools (Brett et al., 2023).

A.2.2. Anatomical registration/segmentation

Prior to registration, both the T1w and T2w images were bias field corrected using *N4BiasFieldCorrection* (Tustison et al., 2010) (spline distance = 100, initial mesh resolution = 1x1x1). Anatomical alignment was accomplished by extracting edge images from each participants' T1w and T2w images using AFNI's 3dedge utility and aligned with a rigid body transform using FLIRT and a correlation ratio cost function. The resulting transformation parameters were saved for later use.

To improve the reliability of the skullstripping procedure, an intermediate anatomical image was utilized. This image was created using the following procedure: (1) T1w and T2w images were rescaled so that values were between 0 and 1; (2) the intermediate volume was calculated where the value at each voxel is determined by the square root of the sum of squares of the corresponding voxels in the scaled T1w and T2w images. The resulting image was passed into BET (Smith, 2002) (fractional intensity threshold = 0.1) for skullstripping. The skullstripping mask was then applied to the participant's aligned T1w and T2w images.

For each participant, we aligned their debiased, skullstripped T1w volume to a common template, the TRIO_Y_NDC atlas, using a 9-parameter affine transformation estimated with FLIRT and a mutual information cost function. We then combined the resulting transformations to align both T1w and T2w volumes to the TRIO atlas using windowed sinc-function interpolation. Anatomical segmentation was done using recon-all from FreeSurfer 6.0 with the -T2 flag enabled (Fischl, 2012).

TRIO_Y_NDC is an atlas based on the 711-2B atlas space. Similar to MNI152 atlas space, a minor 12-parameter affine transform is required to connect the two spaces. This affine transformation mainly reflects two key differences: 1) the 711-2B space is about 5% smaller in linear measure; and 2) it is rotated 3° nose-down about the ear-ear axis. The original atlas-representative image was constructed in 1995 from 12 co-registered T1-weighted images acquired in normal young adults. This image was "spatially normalized" according to the method described by Lancaster et al (Lancaster et al., 1995; Ojemann et al., 1997). Multiple atlas representative images in 711-2B space, e.g., TRIO_Y_NDC, have since been constructed from T1- and T2-weighted images acquired on a variety of scanners.

A.2.3. Alignment of EPI images

Framewise alignment of EPI images was accomplished using AFNI's *3dAllineate*. First, a reference image was constructed by averaging 100 aligned frames from the resting state times series taken from time intervals with the lowest DVARs values (Power et al., 2012). To determine these intervals, a DVARs time series was calculated by computing the temporal derivative of each voxel's time series and normalizing its standard deviation to be equal to 1. For each frame, a DVARs value was calculated as the spatial mean of the absolute value of the derivative time series. Because we were interested in extracting the frames to construct the reference EPI image from prolonged intervals of relatively low DVARs values, we then low-pass filtered the DVARs time series by bi-directionally filtering it with a 1st order digital Butterworth filter with normalized cut-off frequency of 0.16. To create the reference image, we selected 100 frames associated with the lowest smoothed DVARs values and aligned them to the individual frame that had the lowest DVARs value in the session using rigid body transformations. Alignment parameters were estimated with *3dAllineate* and employed a least-squares cost function along with the -autoweight option. Finally we constructed a reference image for the session by averaging the 100 aligned frames together. The reference image was then saved for later use as a target EPI image for *Synth*.

Lastly, framewise alignment parameters were estimated by registering each frame individually to the reference image. To reduce the influence of signal drift on framewise registration and to increase the contribution of properly aligned edges to the cost function, we spatially high-pass filtered the reference functional image and each functional frame independently by smoothing them with large 12 mm FWHM gaussian kernel and subtracting result from the original image. These high-pass filtered volumes were then rigid body aligned using *3dAllineate* which optimized the cost function calculated over 75% of the voxels.

A.2.4. Resting state denoising

Resting state data was denoised by regressing a set of nuisance covariates from each voxel's time series. The design matrix included a constant term; a linear trend component; the first 8 low order Legendre polynomials — to remove low frequency fluctuations; time series corresponding to the 6 framewise rigid-body alignment parameters; and the first derivative of each of the motion parameters along with two additional lagged copies. In addition, the design matrix included nuisance white matter time series as well as time series from nearby voxels outside of the brain used to model global and regional nuisance variability (Petersen et al., 1998; Behzadi et al., 2007; King et al., 2021). The white matter regressors were constructed by extracting and normalizing time series from all voxels labeled as white matter by the participant's FreeSurfer parcellation. Singular value decomposition was then applied to extract 12 time series accounting for the most white matter variance. A similar procedure was used to construct the external signal nuisance regressors. Here, time series from voxels residing within a 5 voxel shell surrounding the brain were extracted. Again, the first 12 singular vectors were computed and included in the denoising design matrix. The complete set of nuisance variables was then projected from the voxel time series using linear regression. Lastly, the regression model residuals were bandpass filtered to retain only frequencies between 0.001 and 0.1 Hz. Frame censoring to remove high motion frames was employed using a weighted correlation approach (Smyser et al., 2010; Power et al., 2012, 2014). First, we

computed a DVARs time series across all brain matter voxels. Frames in which DVARs exceeded two standard deviations were flagged and given a zero weighting. To reduce the influence of any contamination remaining in adjacent frames, the weighting of the three frames preceding and following were also reduced in proportion to their absolute distance from the censored frame.

A.3. Description of radial basis function model used to estimate the synthetic EPI image

Synth allows the user to specify a regression model that maps source image intensities to target image intensities. The radial basis components were constructed using the following procedure: Let $\mathbf{I}_m$ refer to an individual source image, e.g.,T1w or T2w; and $\{c_1, c_2, \ldots, c_J \in C\}$ refers to a set of J evenly spaced values that lie between the maximum and minimum values of $\mathbf{I}_m$. Then each radial basis component, $\mathbf{r}_j$, of an image, $\mathbf{I}_m$, is constructed according to the following equation:

$$\mathbf{r}_j(\mathbf{I}_m) = \exp\left(4\ln(0.5)\left(\frac{\mathbf{I}_m - c_j}{c_1 - c_0}\right)\right) \tag{4}$$

A visual representation of the images produced by applying Eq. (4) to T1w and T2w images is depicted in Fig. 1a. The final design matrix fit by *Synth* consists of each radial basis component concatenated into a single matrix $\mathbf{F}$, along with intercept terms, $\mathbf{e}_m$, (a vector of all ones) for each source image, such that:

$$\mathbf{F} = \begin{bmatrix} | & | & | & | & \cdots & | & \cdots & | \\ \mathbf{e}_0 & \mathbf{e}_1 & \mathbf{r}_0(\mathbf{I}_0) & \mathbf{r}_1(\mathbf{I}_0) & \ldots & \mathbf{r}_0(\mathbf{I}_1) & \ldots & \mathbf{r}_j(\mathbf{I}_i) \\ | & | & | & | & \cdots & | & \cdots & | \end{bmatrix} \tag{5}$$

For both datasets, $\mathbf{F}$ was constructed from a 12 component decomposition on each of the participants' T1w/T2w images, along with all T1w/T2w pairwise interaction terms. As in traditional linear models, these interaction terms are simply the pairwise product of the columns of $\mathbf{F}$.

A.4. The blur operator

While *Synth* estimates the field map corrections at the 1 mm isotropic voxel resolution of the T1w/T2w and synthetic EPI images, the true EPI images are acquired at much lower resolutions. In the case of the MSC dataset, this corresponds to a 4 mm native voxel resolution resampled to 1 mm isotropic resolution; and for the ABCD dataset this corresponds to 2.6 mm native resolution resampled to 1 mm. When lower resolution EPI images are upsampled to the higher resolution grid, their lower resolution manifests as apparent blur. *Synth* models this blur ($\mathbf{E}$ in Eq. (1)) as the convolution of a source image with an Epanachnikov kernel, whose bandwidth is related to the resolution difference between the EPI and anatomical images. This kernel is given by the following expression:

$$\mathbf{E}(\mathbf{x}, h) = \max\left(0, 1 - \frac{1}{h}\|\mathbf{x}\|_2^2\right) \tag{6}$$

where $\mathbf{x} \in \mathbb{R}^3$ is a location on the kernel, and h is the bandwidth parameter. For the results in this paper, the bandwidths of the Epanechnikov kernels were set to $h = 6$ and $h = 10$ for the ABCD and MSC dataset respectively.

Synth models the difference in effective spatial resolution between T1w/T2w and EPI images by blurring the radial basis function (RBF) components of the anatomical images with an Epanachnikov smoothing kernel prior to fitting the model to the target EPI image. In principle, the optical blurring should occur only after the field map deformation is applied, rather than prior to it. However, implementing the model in this way would require repeatedly blurring the high spatial resolution radial basis components of the T1w/T2w source images after every iterative update of the estimated field map parameters. A naive

implementation of this approach would be computationally prohibitive. Prior research has shown that in practice the sort of approximation implemented in *Synth* in many cases does not introduce a significant amount of error (Zhang et al., 2012). However, while the quality of the field map corrections produced by *Synth* appears to be quite high, some improvement may yet be gleaned by using the "correct model", assuming the technical challenges to its implementation can be surmounted.

A.5. Contrast correction for error induced image bias

When source and target images are acquired at substantially different spatial resolutions, or in the early phases of alignment when source and target images are far out of register, the synthetic images produced by *Synth* exhibit significantly "flattened" contrast, even if the rank order voxel intensities comprising the synthetic image largely correspond to those observed in the target image. This is due to the fact that errors in registration or large differences between source and target image resolution bias estimates of the radial basis function parameters, θ, toward the target image mean. In an extreme, but illustrative, case where image registration is so far off that there is no meaningful relationship between corresponding source and target image voxels, the synthetic image will simply be a volume in which all voxels have the same intensity as the target image mean. That is, it will have perfectly 'flat' contrast. As registration improves some, the synthetic image will begin to exhibit more accurately the contrast properties of the target image. In general, however, the contrast of naively estimated synthetic images will tend to be somewhat flattened due to error. To address this, *Synth* includes a tone curve adjustment option that increases contrast of the final synthetic image. This is accomplished by passing the synthetic image voxel intensities through a monotonic non-linear curve, a process similar to tone curve adjustment in traditional image processing. The curve is constrained to reside within the set of beta function cumulative distributions –distributions well suited for modeling tone curves that tend to be sigmoidal– and its parameters, α and β, are optimized by minimizing the negative linear correlation coefficient between the target and synthetic images.

A.6. Estimating field maps with an undistorted synthetic EPI image

The warp parameters that map the undistorted synthetic EPI image to the reference EPI image were estimated using a three step iterative procedure. The final iteration of this procedure produced the estimated field map corrections used for our analyses. This procedure was conducted as follows:

1. Using *Synth*, construct an initial synthetic EPI image based on the current best affine alignment between T1w, T2w and a reference EPI volume. An EPI signal-to-noise mask was used to reduce the contribution of low signal-to-noise areas when constructing the synthetic EPI (see Appendix A.7).
2. Estimate a warp that maps the current synthetic EPI to the reference EPI image. With the exception of the first iteration, each warp estimate is initialized from the warp estimate of the previous iteration.
3. Apply the inverse of this warp, which approximates a field map correction, to the reference EPI image.
4. Repeat steps 1–3, at each iteration, using *Synth* to re-estimate the synthetic EPI image using the most recently corrected EPI image (see Eq. (1)).

For the results presented here, we estimated the field maps using the ANTs SyN algorithm. Each subsequent iteration used a smaller gradient step size, 3, 1, 0.1 for each iteration respectively. SyN parameters were set to use a neighborhood cross-correlation cost function (radius = 2), update and total field variance parameters of 0, winsorize limits

of [0.005, 0.995], histogram matching, BSpline interpolation of order 5 (Lehmann et al., 1999), convergence parameters of 500x0, 1e-6 and 15, smoothing parameters of 0x0, and shrink factor parameters 2x1. In order to leverage the detailed spatial information available in the high resolution synthetic EPI images during the estimation of the field map, the warps at each iteration were constructed in the 1 mm x 1 mm x 1 mm native resolution of the T1w/T2w source images.

A.7. Creating the EPI signal-to-noise mask

To construct a binary mask that would serve to label EPI voxels suffering from a high degree of noise or signal dropout we implemented a linear discriminant analysis procedure. We began by using the brain mask derived from the skullstripping stage of the anatomical volume and aligned it to the functional dataset using the affine transformation generated from *Synth*. This mask served as an initial guess, or prior estimate, for which voxels were likely to represent meaningful fMRI signals in the resting state EPI dataset, and those which represent noise and areas of signal dropout Supplemental Figure 3. Using this initial labeling, we implemented a between-class linear discriminant analysis (LDA) based, first, on the raw un-centered and un-scaled voxel time series and then computed the projection of each voxel's time series onto the resulting principal discriminant vector. We repeated this LDA procedure on the centered and normalized time series from each voxel, again projecting the times series back onto the principal discriminant vector. Thus, each voxel was associated with two projection values. We then created a third value associated with each voxel by computing the product of the two principal discriminant vector projections. A second level of LDA was done on these three values. Lastly, Otsu's method (Otsu, 1979) was used to label each voxel as 'signal' or 'noise', depending upon its projection onto this final 3-element principal discriminant vector. We refined this labeling by repeating the entire LDA procedure over 20 iterations. During each subsequent iteration the labels produced by the previous iteration were used during the LDA procedure. Sufficiently iterated, this approach converges to a stable binary signal-or-noise label for each voxel. Lastly, to construct a mask containing only within-brain noise voxels –voxels most likely to be affected by signal dropout– we masked the final result of this iterative process to remove non-brain voxels (defined by the skullstrip mask) and morphologically open and dilate the image by 1 voxel. This mask was used as a weighting volume to further refine the estimate of the synthetic functional EPI by reducing the contributions of brain regions where no correspondence could exist between EPI images and T1w/T2w images due to signal drop out.

A.8. Distortion correction using alternating descent optimization with synth

Distortion correction with *Synth* is accomplished through an alternating descent optimization approach, which solves for each parameter individually while keeping the others fixed. At the end of each iteration of alternating descent, updates to the distortion correction estimation improves the accuracy of the synthetic image constructed during the subsequent iteration. These sequential improvements allow for iteratively more accurate distortion correction. The complete procedure is described in Algorithm 1.

Prior to the alternating descent procedure, initial warp parameters, ϕ_0, are initialized by a rigid-body transform that aligns the T2w anatomical image and EPI image using a mutual information cost function. The rigid-body parameters from this step are then converted into a displacement field and applied to the EPI image, aligning it to the anatomical images that underlie the matrix, $\mathbf{F}$. Algorithm 1 begins by constructing the radial basis function matrix, $\mathbf{F}$, from the T1w/T2w anatomical images. Then the main alternating minimization loop is run for each iteration n. First, θ_n is computed by solving the linear system, $f^{-1}\left(\mathbf{y};\phi_{n-1}\right) = \mathbf{EF}\theta_n$. Here, f^{-1} denotes the inverse displacement field that non-linearly transforms the real EPI image to the space of the synthetic EPI image. θ_n is then used to construct an intermediate synthetic image. Next, global contrast parameters, α_n and β_n are estimated to improve overall similarity between the EPI image and the intermediate synthetic image. Finally, a new estimate for the displacement field parameters, ϕ_n, are estimated using the ANTs SyN non-linear warping algorithm. Each iteration updates ϕ_n, θ_n, α_n, and β_n until a specified number of iterations, N, is reached.

Appendix B. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.dcn.2023.101234.

References

Andersson, J.L.R., Skare, S., 2002. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. Neuroimage 16 (1), 177–199.

Andersson, J.L.R., Skare, S., Ashburner, J., 2003. How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. Neuroimage 20 (2), 870–888.

Andersson, J.L.R., Sotiropoulos, S.N., 2016. An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. Neuroimage 125, 1063–1078.

Antun, V., Renna, F., Poon, C., Adcock, B., Hansen, A.C., 2020. On instabilities of deep learning in image reconstruction and the potential costs of AI. Proc. Natl. Acad. Sci. USA 117 (48), 30088–30095.

Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C., 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Med. Image Anal. 12 (1), 26–41.

Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C., 2011. A reproducible evaluation of ANTs similarity metric performance in brain image registration. Neuroimage 54 (3), 2033–2044.

Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A.Z., Van Essen, D.C., WU-Minn HCP Consortium, 2013. Function in the human connectome: task-fMRI and individual differences in behavior. Neuroimage 80, 169–189.

Behzadi, Y., Restom, K., Liau, J., Liu, T.T., 2007. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. Neuroimage 37 (1), 90–101.

Bhadra, S., Kelkar, V.A., Brooks, F.J., Anastasio, M.A., 2021. On hallucinations in tomographic image reconstruction. IEEE Trans. Med. Imaging 40 (11), 3249–3260.

Braga, R.M., Buckner, R.L., 2017. Parallel interdigitated distributed networks within the individual estimated by intrinsic functional connectivity. Neuron 95 (2), 457–471.e5.

Brett, M., 2023. transforms3d. URL https://github.com/matthew-brett/transforms3d.

Brett, M., Markiewicz, C.J., Hanke, M., Côté, M.-A., Cipollini, B., McCarthy, P., Jarecka, D., Cheng, C.P., Halchenko, Y.O., Cottaar, M., Larson, E., Ghosh, S., Wassermann, D., Gerhard, S., Lee, G.R., Wang, H.-T., Kastman, E., Kaczmarzyk, J., Guidotti, R., Daniel, J., Duek, O., Rokem, A., Madison, C., Papadopoulos Or-fanos, D., Sólon, A., Moloney, B., Morency, F.C., Goncalves, M., Baratz, Z., Markello, R., Riddell, C., Burns, C., Millman, J., Gramfort, A., Leppäkangas, J., van den Bosch, J.J., Vincent, R.D., Braun, H., Subramaniam, K., Van, A., Gor-golewski, K.J., Raamana, P.R., Klug, J., Nichols, B.N., Baker, E.M., Hayashi, S., Pinsard, B., Haselgrove, C., Hymers, M., Esteban, O., Koudoro, S., Pérez-García, F., Dockès, J., Oosterhof, N.N., Amirbekian, B., Nimmo-Smith, I., Nguyen, L., Reddi-gari, S., St-Jean, S., Panfilov, E., Garyfallidis, E., Varoquaux, G., Legarreta, J.H., Hahn, K.S., Waller, L., Hinds, O.P., Fauber, B., Roberts, J., Poline, J.-B., Stutters, J., Jordan, K., Cieslak, M., Moreno, M.E., Hrnoiar, T., Haenel, V., Schwartz, Y., Dar-win, B.C., Thirion, B., Gauthier, C., Solovey, I., Gonzalez, I., Palasubramaniam, J., Lecher, J., Leinweber, K., Raktivan, K., Calábková, M., Fischer, P., Gervais, P., Gadde, S., Ballinger, T., Roos, T., Reddam, V.R., freec84, 2023. Nipy/nibabel: 5.0.1. http://dx.doi.org/10.5281/zenodo.7633628.

Callaghan, P.T., 1990. Susceptibility-limited resolution in nuclear magnetic resonance microscopy. J. Magn. Reson. 87 (2), 304–318.

Casey, B.J., Cannonier, T., Conley, M.I., Cohen, A.O., Barch, D.M., Heitzeg, M.M., Soules, M.E., Teslovich, T., Dellarco, D.V., Garavan, H., Orr, C.A., Wager, T.D., Banich, M.T., Speer, N.K., Sutherland, M.T., Riedel, M.C., Dick, A.S., Bjork, J.M., Thomas, K.M., Chaarani, B., Mejia, M.H., Hagler, Jr., D.J., Daniela Cornejo, M., Sicat, C.S., Harms, M.P., Dosenbach, N.U.F., Rosenberg, M., Earl, E., Bartsch, H., Watts, R., Polimeni, J.R., Kuperman, J.M., Fair, D.A., Dale, A.M., ABCD Imaging Acquisition Workgroup, 2018. The adolescent brain cognitive development (ABCD) study: Imaging acquisition across 21 sites. Dev. Cogn. Neurosci. 32, 43–54.

Chambers, M.C., Bhushan, C., Haldar, J.P., Leahy, R.M., Shattuck, D.W., 2015. Correct-ing inhomogeneity-induced distortion in FMRI using non-rigid registration. In: 2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI). pp. 1364–1367.

Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput. Biomed. Res. 29 (3), 162–173.

Cusack, R., Brett, M., Osswald, K., 2003. An evaluation of the use of magnetic field maps to undistort echo-planar images. Neuroimage 18 (1), 127–142.

Devlin, J.T., Russell, R.P., Davis, M.H., Price, C.J., Wilson, J., Moss, H.E., Matthews, P.M., Tyler, L.K., 2000. Susceptibility-induced loss of signal: comparing PET and fMRI on a semantic task. Neuroimage 11 (6 Pt 1), 589–600.

Dickie, E.W., Anticevic, A., Smith, D.E., Coalson, T.S., Manogaran, M., Calarco, N., Viviano, J.D., Glasser, M.F., Van Essen, D.C., Voineskos, A.N., 2019. Ciftify: A framework for surface-based analysis of legacy MR acquisitions. Neuroimage 197, 818–826.

Esteban, O., Adebimpe, A., Markiewicz, C.J., Goncalves, M., Blair, R.W., Cieslak, P., Naveau, M., Sitek, K.R., Sneve, M.H., Provins, C., al., E., 2021. The Bermuda Triangle of d- and f-MRI sailors - software for susceptibility distortions (SDCFlows).

Esteban, O., Markiewicz, C.J., Blair, R.W., Moodie, C.A., Isik, A.I., Erramuzpe, A., Kent, J.D., Goncalves, M., DuPre, E., Snyder, M., Oya, H., Ghosh, S.S., Wright, J., Durnez, J., Poldrack, R.A., Gorgolewski, K.J., 2019. fMRIPrep: a robust preprocessing pipeline for functional MRI. Nature Methods 16 (1), 111–116.

Fischl, B., 2012. FreeSurfer. Neuroimage 62 (2), 774–781.

Gholipour, A., Kehtarnavaz, N., Briggs, R.W., Gopinath, K.S., Ringe, W., Whittemore, A., Cheshkov, S., Bakhadirov, K., 2008a. Validation of non-rigid registration between functional and anatomical magnetic resonance brain images. IEEE Trans. Biomed. Eng. 55 (2 Pt 1), 563–571.

Gholipour, A., Kehtarnavaz, N., Gopinath, K., Briggs, R., 2008b. Cross-Validation of deformable registration with field maps in functional magnetic resonance brain imaging. IEEE J. Sel. Top. Signal Process. 2 (6), 854–869.

Glasser, M.F., Smith, S.M., Marcus, D.S., Andersson, J.L.R., Auerbach, E.J., Behrens, T.E.J., Coalson, T.S., Harms, M.P., Jenkinson, M., Moeller, S., Robinson, E.C., Sotiropoulos, S.N., Xu, J., Yacoub, E., Ugurbil, K., Van Essen, D.C., 2016. The Human Connectome Project's neuroimaging approach. Nature Neurosci. 19 (9), 1175–1187.

Glover, G.H., 2012. Spiral imaging in fMRI. Neuroimage 62 (2), 706–712.

Gordon, E.M., Laumann, T.O., Gilmore, A.W., Newbold, D.J., Greene, D.J., Berg, J.J., Ortega, M., Hoyt-Drazen, C., Gratton, C., Sun, H., Hampton, J.M., Coalson, R.S., Nguyen, A.L., McDermott, K.B., Shimony, J.S., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., Nelson, S.M., Dosenbach, N.U.F., 2017. Precision functional mapping of individual human brains. Neuron 95 (4), 791–807.e7.

Gordon, E.M., Laumann, T.O., Marek, S., Raut, R.V., Gratton, C., Newbold, D.J., Greene, D.J., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., Dosenbach, N.U.F., Nelson, S.M., 2020. Default-mode network streams for coupling to language and control systems. Proc. Natl. Acad. Sci. USA 117 (29), 17308–17319.

Gordon, E.M., Lynch, C.J., Gratton, C., Laumann, T.O., Gilmore, A.W., Greene, D.J., Ortega, M., Nguyen, A.L., Schlaggar, B.L., Petersen, S.E., Dosenbach, N.U.F., Nelson, S.M., 2018. Three distinct sets of connector hubs integrate human brain function. Cell Rep. 24 (7), 1687–1695.e4.

Gorgolewski, K.J., Auer, T., Calhoun, V.D., Craddock, R.C., Das, S., Duff, E.P., Flandin, G., Ghosh, S.S., Glatard, T., Halchenko, Y.O., et al., 2016. The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. Sci. Data 3 (1), 1–9.

Graham, M.S., Drobnjak, I., Jenkinson, M., Zhang, H., 2017. Quantitative assessment of the susceptibility artefact and its interaction with motion in diffusion MRI. PLoS One 12 (10), e0185647.

Gratton, C., Laumann, T.O., Nielsen, A.N., Greene, D.J., Gordon, E.M., Gilmore, A.W., Nelson, S.M., Coalson, R.S., Snyder, A.Z., Schlaggar, B.L., Dosenbach, N.U.F., Petersen, S.E., 2018. Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. Neuron 98 (2), 439–452.e5.

Greene, D.J., Black, K.J., Schlaggar, B.L., 2016. Considerations for MRI study design and implementation in pediatric and clinical populations. Dev. Cognit. Neurosci. 18, 101–112.

Greene, D.J., Marek, S., Gordon, E.M., Siegel, J.S., Gratton, C., Laumann, T.O., Gilmore, A.W., Berg, J.J., Nguyen, A.L., Dierker, D., Van, A.N., Ortega, M., Newbold, D.J., Hampton, J.M., Nielsen, A.N., McDermott, K.B., Roland, J.L., Norris, S.A., Nelson, S.M., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., Dosenbach, N.U.F., 2020. Integrative and Network-Specific connectivity of the basal ganglia and thalamus defined in individuals. Neuron 105 (4), 742–758.e6.

Greve, D.N., Fischl, B., 2009. Accurate and robust brain image alignment using boundary-based registration. Neuroimage 48 (1), 63–72.

Gureyev, T.E., Kozlov, A., Paganin, D.M., Nesterets, Y.I., Quiney, H.M., 2020. Noise-resolution uncertainty principle in classical and quantum systems. Sci. Rep. 10 (1), 7890.

Hong, X., To, X.V., Teh, I., Soh, J.R., Chuang, K.-H., 2015. Evaluation of EPI distortion correction methods for quantitative MRI of the brain at high magnetic field. Magn. Reson. Imaging 33 (9), 1098–1105.

Huntenburg, J.M., 2014. (Ph.D. thesis). Freie Universität Berlin.

Jernigan, T.L., Brown, S.A., Dowling, G.J., 2018. The adolescent brain cognitive development study. J. Res. Adolesc. 28 (1), 154–156.

Jezzard, P., Balaban, R.S., 1995. Correction for geometric distortion in echo planar images from B0 field variations. Magn. Reson. Med. 34 (1), 65–73.

Jezzard, P., Barnett, A.S., Pierpaoli, C., 1998. Characterization of and correction for eddy current artifacts in echo planar diffusion imaging. Magn. Reson. Med. 39 (5), 801–812.

King, L.S., Camacho, M.C., Montez, D.F., Humphreys, K.L., Gotlib, I.H., 2021. Naturalistic language input is associated with Resting-State functional connectivity in infancy. J. Neurosci. 41 (3), 424–434.

Klein, A., Andersson, J., Ardekani, B.A., Ashburner, J., Avants, B., Chiang, M.-C., Christensen, G.E., Collins, D.L., Gee, J., Hellier, P., Song, J.H., Jenkinson, M., Lepage, C., Rueckert, D., Thompson, P., Vercauteren, T., Woods, R.P., Mann, J.J., Parsey, R.V., 2009. Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. Neuroimage 46 (3), 786–802.

Kybic, J., Thevenaz, P., Nirkko, A., Unser, M., 2000. Unwarping of unidirectionally distorted EPI images. IEEE Trans. Med. Imaging 19 (2), 80–93.

Lancaster, J.L., Glass, T.G., Lankipalli, B.R., Downs, H., Mayberg, H., Fox, P.T., 1995. A modality-independent approach to spatial normalization of tomographic images of the human brain. Hum. Brain Mapping 3 (3), 209–223.

Laumann, T.O., Gordon, E.M., Adeyemo, B., Snyder, A.Z., Joo, S.J., Chen, M.-Y., Gilmore, A.W., McDermott, K.B., Nelson, S.M., Dosenbach, N.U.F., Schlaggar, B.L., Mumford, J.A., Poldrack, R.A., Petersen, S.E., 2015. Functional system and areal organization of a highly sampled individual human brain. Neuron 87 (3), 657–670.

Lehmann, T.M., Gönner, C., Spitzer, K., 1999. Survey: interpolation methods in medical image processing. IEEE Trans. Med. Imaging 18 (11), 1049–1075.

Lynch, C.J., Power, J.D., Scult, M.A., Dubin, M., Gunning, F.M., Liston, C., 2020. Rapid precision functional mapping of individuals using Multi-Echo fMRI. Cell Rep. 33 (12), 108540.

Marek, S., Siegel, J.S., Gordon, E.M., Raut, R.V., Gratton, C., Newbold, D.J., Ortega, M., Laumann, T.O., Adeyemo, B., Miller, D.B., Zheng, A., Lopez, K.C., Berg, J.J., Coalson, R.S., Nguyen, A.L., Dierker, D., Van, A.N., Hoyt, C.R., McDermott, K.B., Norris, S.A., Shimony, J.S., Snyder, A.Z., Nelson, S.M., Barch, D.M., Schlaggar, B.L., Raichle, M.E., Petersen, S.E., Greene, D.J., Dosenbach, N.U.F., 2018. Spatial and temporal organization of the individual human cerebellum. Neuron 100 (4), 977–993.e7.

Marek, S., Tervo-Clemmens, B., Nielsen, A.N., Wheelock, M.D., Miller, R.L., Laumann, T.O., Earl, E., Foran, W.W., Cordova, M., Doyle, O., Perrone, A., Miranda-Dominguez, O., Feczko, E., Sturgeon, D., Graham, A., Hermosillo, R., Snider, K., Galassi, A., Nagel, B.J., Ewing, S.W.F., Eggebrecht, A.T., Garavan, H., Dale, A.M., Greene, D.J., Barch, D.M., Fair, D.A., Luna, B., Dosenbach, N.U.F., 2019. Identifying reproducible individual differences in childhood functional brain networks: An ABCD study. Dev. Cogn. Neurosci. 40, 100706.

Newbold, D.J., Gordon, E.M., Laumann, T.O., Montez, D.F., et al., 2020a. Cingulo-Opercular control network supports disused motor circuits in standby mode. BioRxiv.

Newbold, D.J., Laumann, T.O., Hoyt, C.R., Hampton, J.M., Montez, D.F., Raut, R.V., Ortega, M., Mitra, A., Nielsen, A.N., Miller, D.B., Adeyemo, B., Nguyen, A.L., Scheidter, K.M., Tanenbaum, A.B., Van, A.N., Marek, S., Schlaggar, B.L., Carter, A.R., Greene, D.J., Gordon, E.M., Raichle, M.E., Petersen, S.E., Snyder, A.Z., Dosenbach, N.U.F., 2020b. Plasticity and spontaneous activity pulses in disused human brain circuits. Neuron 107 (3), 580–589.e6.

Ojemann, J.G., Akbudak, E., Snyder, A.Z., McKinstry, R.C., Raichle, M.E., Conturo, T.E., 1997. Anatomic localization and quantitative analysis of gradient refocused echo-planar fMRI susceptibility artifacts. Neuroimage 6 (3), 156–167.

Otsu, N., 1979. A threshold selection method from Gray-Level histograms. IEEE Trans. Syst. Man Cybern. 9 (1), 62–66.

Paul, M., 2021. FNIRT/UserGuide - FslWiki. https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/FNIRT/UserGuide, (Accessed 12 February 2021).

Petersen, N.V., Jensen, J.L., Burchhardt, J., Stø dkilde Jø rgensen, H., 1998. State space models for physiological noise in fMRI time series. Neuroimage 7 (4), S592.

Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage 59 (3), 2142–2154.

Power, J.D., Mitra, A., Laumann, T.O., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2014. Methods to detect, characterize, and remove motion artifact in resting state fMRI. Neuroimage 84, 320–341.

Rohlfing, T., 2012. Image similarity and tissue overlaps as surrogates for image registration accuracy: Widely used but unreliable. IEEE Trans. Med. Imaging 31 (2), 153–163.

Saad, Z.S., Glen, D.R., Chen, G., Beauchamp, M.S., Desai, R., Cox, R.W., 2009. A new method for improving functional-to-structural MRI alignment using local Pearson correlation. Neuroimage 44 (3), 839–848.

Schallmo, M.-p., Weldon, K.B., Burton, P.C., Sponheim, S.R., Olman, C.A., 2021. Assessing methods for geometric distortion compensation in 7 T gradient echo functional MRI data. 42, (13), pp. 4205–4223,

Schilling, K.G., Blaber, J., Hansen, C., Cai, L., Rogers, B., Anderson, A.W., Smith, S., Kanakaraj, P., Rex, T., Resnick, S.M., Shafer, A.T., Cutting, L.E., Woodward, N., Zald, D., Landman, B.A., 2020. Distortion correction of diffusion weighted MRI without reverse phase-encoding scans or field-maps. PLoS One 15 (7), e0236418.

Smith, S.M., 2002. Fast robust automated brain extraction. Hum. Brain Mapp. 17 (3), 143–155.

Smyser, C.D., Inder, T.E., Shimony, J.S., Hill, J.E., Degnan, A.J., Snyder, A.Z., Neil, J.J., 2010. Longitudinal analysis of neural network development in preterm infants. Cereb. Cortex 20 (12), 2852–2862.

Sotiras, A., Davatzikos, C., Paragios, N., 2013. Deformable medical image registration: a survey. IEEE Trans. Med. Imaging 32 (7), 1153–1190.

Studholme, C., Constable, R.T., Duncan, J.S., 2000. Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion model. IEEE Trans. Med. Imaging 19 (11), 1115–1127.

Sylvester, C.M., Yu, Q., Srivastava, A.B., Marek, S., Zheng, A., Alexopoulos, D., Smyser, C.D., Shimony, J.S., Ortega, M., Dierker, D.L., Patel, G.H., Nelson, S.M., Gilmore, A.W., McDermott, K.B., Berg, J.J., Drysdale, A.T., Perino, M.T., Snyder, A.Z., Raut, R.V., Laumann, T.O., Gordon, E.M., Barch, D.M., Rogers, C.E., Greene, D.J., Raichle, M.E., Dosenbach, N.U.F., 2020. Individual-specific functional connectivity of the amygdala: A substrate for precision psychiatry. Proc. Natl. Acad. Sci. USA 117 (7), 3808–3818.

Togo, H., Rokicki, J., Yoshinaga, K., Hisatsune, T., Matsuda, H., Haga, N., Hanakawa, T., 2017. Effects of Field-Map distortion correction on resting state functional connectivity MRI. Front. Neurosci. 11, 656.

Treiber, J.M., White, N.S., Steed, T.C., Bartsch, H., Holland, D., Farid, N., McDonald, C.R., Carter, B.S., Dale, A.M., Chen, C.C., 2016. Characterization and correction of geometric distortions in 814 diffusion weighted images. PLoS One 11 (3), e0152472.

Tustison, N.J., Avants, B.B., Cook, P.A., Zheng, Y., Egan, A., Yushkevich, P.A., Gee, J.C., 2010. N4ITK: Improved N3 bias correction. IEEE Trans. Med. Imaging 29 (6), 1310–1320.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium, 2013. The WU-Minn Human Connectome Project: an overview. Neuroimage 80, 62–79.

Visser, M., Embleton, K.V., Jefferies, E., Parker, G.J., Ralph, M.A.L., 2010a. The inferior, anterior temporal lobes and semantic memory clarified: novel evidence from distortion-corrected fMRI. Neuropsychologia 48 (6), 1689–1696.

Visser, E., Qin, S., Zwiers, M., 2010b. EPI distortion correction by constrained nonlinear coregistration improves group fMRI. In: Proc. 18th Ann. Mtg. of ISMRM, Vol. 18. p. 3459.

Wang, S., Peterson, D.J., Gatenby, J.C., Li, W., Grabowski, T.J., Madhyastha, T.M., 2017. Evaluation of field map and nonlinear registration methods for correction of susceptibility artifacts in diffusion MRI. Front. Neuroinform. 11, 17.

West, J.B., Michael Fitzpatrick, J., Wang, M.Y., Dawant, B.M., Maurer, Jr., C.R., Kessler, R.M., Maciunas, R.J., Barillot, C., Lemoine, D., Collignon, A.M.F., Maes, F., Suetens, P., Vandermeulen, D., van den Elsen, P.A., Hemler, P.F., Napel, S., Sumanaweera, T.S., Harkness, B.A., Hill, D.L., Studholme, C., Malandain, G., Pennec, X., Noz, M.E., Maguire, Jr., G.Q., Pollack, M., Pelizzari, C.A., Robb, R.A., Hanson, D.P., Woods, R.P., 1996. Comparison and evaluation of retrospective intermodality image registration techniques. Medical Imaging 1996: Image Processing.

Zhang, X., Jiang, J., Peng, S., 2012. Commutability of blur and affine warping in super-resolution with application to joint estimation of triple-coupled variables. IEEE Trans. Image Process. 21 (4), 1796–1808.

Zhark the Grotesquely Warped (but still strangely handsome), 2021. AFNI program: 3dQwarp. https://afni.nimh.nih.gov/pub/dist/doc/program_help/3dQwarp.html. (Accessed 16 February 2021).

Zheng, A., Montez, D.F., Marek, S., Gilmore, A.W., Newbold, D.J., Laumann, T.O., Kay, B.P., Seider, N.A., Van, A.N., Hampton, J.M., Alexopoulous, D., Schlaggar, B.L., Sylvester, C.M., Greene, D.J., Shimony, J.S., Nelson, S.M., Wig, G.S., Gratton, C., McDermott, K.B., Raichle, M.E., Gordon, E.M., Dosenbach, N.U.F., 2020. Parallel Hippocampal-Parietal circuits for self- and goal-oriented processing. Cold Spring Harbor Lab. 2020.12.01.395210.