Gopinath, K., Hoopes, A., Alexander, D.C., Arnold, S.E., Balbastre, Y., Casamitjanam A., Cheng, Y., Chua, R.Y.Z., Edlow, B.L., Fischl, B., Gazula, H., Hoffmann, M., Keene, C.D., Kim, S., Kimberly, W.T., Laguna, S., Larson, K.E., Van Leemput, K., Puonti, O., Rodrigues, L.M., Rosen, M.S., Tregidgo, H.F.J., Varadarajan, D., Young, S.I., Dalca, A.V., & Iglesias, J.E. (2024). Synthetic data in generalizable, learning-based neuroimaging. *Imaging Neuroscience*, Advance Publication. https://doi.org/10.1162/imag_a_00337

Synthetic data in generalizable, learning-based neuroimaging

Karthik Gopinath[*,1], Andrew Hoopes[*,1,2], Daniel C. Alexander[3], Steven E. Arnold[1], Yael Balbastre[1], Adrià Casamitjana[4], You Cheng[1], Russ Yue Zhi Chua[1,2], Brian L. Edlow[1], Bruce Fischl[1], Harshvardhan Gazula[2], Malte Hoffmann[1], C. Dirk Keene[5], Seunghoi Kim[3], W. Taylor Kimberly[1], Sonia Laguna[6], Kathleen E. Larson[1], Koen Van Leemput[1], Oula Puonti[1,7], Livia M. Rodrigues[1,8], Matthew S. Rosen[1], Henry F. J. Tregidgo[3], Divya Varadarajan[1], Sean I. Young[1,2], Adrian V. Dalca[†,1,2], Juan Eugenio Iglesias[†,1,2,3]

[1] Massachusetts General Hospital, Harvard Medical School, Boston, United States

[2] Massachusetts Institute of Technology, Cambridge, United States

[3] University College London, England

[4] Universitat de Girona, Spain

[5] University of Washington, Seattle, United States

[6] ETH Zürich, Switzerland

[7] Copenhagen University Hospital, Denmark

[8] Universidade Estadual de Campinas, São Paulo, Brazil

[*] Shared authorship

[†] Shared authorship

Correspondence: K. Gopinath (kgopinath@mgh.harvard.edu), A. Hoopes (ahoopes@mgh.harvard.edu)

Keywords: SynthSeg, SynthStrip, SynthMorph, EasyReg, SynthSR

Abstract: Synthetic data has emerged as an attractive option for developing machine learning methods in human neuroimaging, particularly in magnetic resonance imaging (MRI) – a modality where image contrast depends enormously on acquisition hardware and parameters. This retrospective paper reviews a family of recently proposed methods, based on synthetic data, for generalizable machine learning in brain MRI analysis. Central to this framework is the concept of domain randomization, which involves training neural networks on a vastly diverse array of synthetically generated images with random contrast properties. This technique has enabled robust, adaptable models that are capable of handling diverse MRI contrasts, resolutions, and pathologies, while working out-of-the-box, without retraining. We have successfully applied this method to tasks such as whole brain segmentation (SynthSeg), skull-stripping (SynthStrip), registration (SynthMorph, EasyReg), super-resolution and MR contrast transfer (SynthSR). Beyond these applications, the paper discusses other possible use cases and future work in our methodology. Neural networks trained with synthetic data enable the analysis of clinical MRI, including large retrospective datasets, while greatly alleviating (and sometimes eliminating) the need for substantial labeled datasets, and offer enormous potential as robust tools to address various research goals.

1 Introduction

As MRI provided researchers with the opportunity to study the human brain *in vivo*, open-source neuroimaging software emerged to enable quantitative analysis of imaging data at scale. Packages like FreeSurfer (Fischl, 2012), FSL (S. M. Smith et al., 2004), SPM (Ashburner & Friston, 2005), or AFNI (Cox, 1996) have facilitated large-scale studies of healthy aging, dementia, and neurological disorders (Dima et al., 2022; Frangou et al., 2022; The Alzheimer's Disease Neuroimaging Initiative et al., 2015; Thompson et al., 2020). Other tools analyze scans of subjects with diseases that more severely alter the structure of the brain, such as strokes or tumors (Gordillo et al., 2013; Kamnitsas et al., 2017; Pereira et al., 2016). These tools have also increased the reproducibility of research, particularly when combined with public datasets, such as ADNI (Jack Jr. et al., 2008), HCP (Van Essen et al., 2013), the UK BioBank (Alfaro-Almagro et al., 2018), or BraTS (Menze et al., 2015).

Over the past decade, rapid advances in deep learning have transformed neuroimaging methods, in areas such as MRI reconstruction (Zhu et al., 2018), segmentation (Kamnitsas et al., 2017; Milletari et al.,

2016), or registration (Balakrishnan et al., 2019; De Vos et al., 2019). Powered by convolutional neural networks (CNNs, LeCun et al. 1998) or, more recently, vision transformers (Z. Liu et al., 2021), these methods frequently achieve higher levels of performance than their classical counterparts. In addition, deep learning methods often have considerably shorter inference run times compared to classical methods, especially when running on graphics processing units (GPUs). This increased speed enables application at a larger scale or in time-sensitive applications, such as clinical fetal imaging (Ebner et al., 2020; Hoffmann, Abaci Turk, et al., 2021).

However, a major roadblock to the wider adoption of deep learning methods in neuroimaging is their sensitivity to the so-called "domain shift" – the drop in performance when models trained on one dataset are applied to other datasets with even slightly different image intensity profiles (Billot, Greve, et al., 2023). The domain shift is particularly problematic in uncalibrated modalities like MRI, as opposed to computerized tomography, where voxel intensities correspond to physical Hounsfield units. In MRI, variations in image intensities arise from differences in hardware, pulse sequences, slice direction, or resolution.

While data augmentation (Shorten & Khoshgoftaar, 2019; A. Zhao et al., 2019) and domain adaptation techniques (Wang & Deng, 2018) mitigate the problem for smaller domain shifts, they do not close the domain gap – especially in more dramatic shifts. For this reason, most algorithms in neuroimaging packages (FreeSurfer, FSL, SPM) still rely on Bayesian methods that are robust against changes to MRI contrast (Ashburner & Friston, 2005; Fischl et al., 2002; Patenaude et al., 2011) or pathology (Agn et al., 2019; Van Leemput et al., 2001).

Recently, our group has proposed a solution to this problem based on "domain randomization" (Tobin et al., 2017; Tremblay et al., 2018). These methods rely on training neural networks with an extremely wide distribution of synthetic data simulated with random parameters, in our case, MRI contrast, resolution, and artifacts like bias field or motion. By randomizing these parameters at each iteration during training, the neural network learns *not* to expect a specific MRI contrast or resolution, thus becoming agnostic to them. With this strategy, real images appear to the trained neural network as just another variation of the wide distribution.

Domain randomization has the advantage that the trained neural networks are available to process new images "out of the box", without retraining or domain adaptation. This feature makes domain randomization particularly appealing when publicly distributing software. In FreeSurfer, we currently distribute multiple domain randomization powered methods for an array of brain MRI analysis tasks, including segmentation (Billot, Greve, et al., 2023; Billot, Magdamo, et al., 2023; Hoopes, Mora, et al., 2022; Kelley et al., 2024), registration (Hoffmann et al., 2022, 2023, 2024; Iglesias, 2023), super-

resolution (Iglesias et al., 2020, 2022), and contrast transfer (Iglesias et al., 2023).

Throughout the rest of this article, we do not make any assumption on the specific architectures used for learning. Instead, we focus on the *data* that are used for training them, since architecture and training data are generally independent of each other. Of course, some architectures are more suitable to certain problems than other; the reader is referred to Alzubaidi et al., 2021; Khan et al., 2023; Taye, 2023 for a comprehensive review of this topic.

2 Background

Machine learning models developed with partially representative datasets or rigid simulation environments often struggle to adapt to the unpredictability and complexity of real-world inputs. Consequently, integrating data synthesis into model training pipelines has become a wide-spread strategy to improve model generalization under such circumstances. In this section, we review existing efforts to generate training data in the context of both general computer vision and medical image analysis.

2.1 Synthetic data in computer vision

Many real-world imaging datasets lack the diversity and scope necessary to train universally robust solutions for computer vision tasks. This is in part because acquiring images with useful annotations – such as semantic segmentations, point-wise labels, or classifications – requires substantial manual effort and, in some cases, costly budgets. Consequently, several groups have developed large-scale, specialized datasets comprising training images and annotations generated through automated means.

Some of these datasets derive from spatial augmentation or corruption pipelines, for instance producing image-annotation pairs for object anomaly and defect analysis (C.-L. Li et al., 2021). Alternatively, recent generative approaches learn to synthesize a distribution of high-quality images, in some cases containing features described by a input language prompt. This facilitates building specialized datasets for object detection entirely from scratch, manipulating images, or approximating missing data in existing cohorts (Azizi et al., 2023; Kar et al., 2019; G. Zhang et al., 2021).

Other synthetic data contributions use techniques in computer graphics to simulate images from 3D object representations. These methods compute ground-truth attributes, like pixel class, depth, and location, used to train models for autonomous vehicles (Alhaija et al., 2017; Das et al., 2023; Kundu et al., 2018; Ros et al., 2016; Sadeghi & Levine, 2016; Wu et al., 2017), human pose estimation (Hewitt et al.,

2023; Varol et al., 2017), or facial analysis (Bae et al., 2023; Wood et al., 2021, 2022). Some of these approaches involve highly complex simulations. For instance, robotics models use physics engines to synthesize multimodal inputs from interactive environments with contact dynamics and constraints (Shah et al., 2018; Todorov et al., 2012). Similarly, synthetic scene generation creates customizable environments for the development and evaluation of autonomous driving models in controlled scenarios (Dosovitskiy et al., 2017). These techniques not only improve model robustness for realistic inference-time objectives, but also provide control over certain data biases and anonymity, especially for privacy-sensitive domains like healthcare or facial image analysis.

2.2 Applications of synthetic data in medical image analysis

Synthesis methods are used to overcome similar data scarcity challenges in medical imaging. These approaches predominantly use synthetic and augmented datasets of diverse human anatomies to balance imaging features or disease classes, reduce manual annotation efforts, and ultimately enhance model robustness for tasks such as segmentation, registration, and classification. Some applications develop structural datasets for anatomies, such as the heart (Al Khalil et al., 2020; Azizmohammadi et al., 2022; Xanthis et al., 2021), kidney (Brumer et al., 2022), vertebrae (Sun et al., 2023), and brain (Billot, Greve, et al., 2023; Dorjsembe et al., 2024; Kossen et al., 2022), or underrepresented populations, such as pediatric data (de Dumast et al., 2022). Synthesis techniques are particularly valuable for obtaining data that are difficult to collect in clinical or research settings. For instance, numerous published datasets with generated lesions, and often annotations, include those that synthesize lung disease and COVID-19 lesions (Gohorbani et al., 2019; Y. Jiang et al., 2020; X. Li et al., 2009; Thambawita et al., 2022; Waheed et al., 2020; Zunair & Hamza, 2021), brain tumors and neurodegenerative conditions (Ahmad et al., 2022; Bernal et al., 2021; Prados et al., 2016; Salem et al., 2019; H.-C. Shin et al., 2018), breast cancer (Badano et al., 2018; Cha et al., 2020; Pezeshk et al., 2015; Sarno et al., 2021; Sauer et al., 2023; Sizikova et al., 2024), colon polyps (Thambawita et al., 2022), and liver lesions (Ben-Cohen et al., 2019; Frid-Adar et al., 2018).

Some applications of medical image synthesis employ style transfer techniques to alter the modality or contrast properties of real scans (J. Jiang et al., 2018; J. Li et al., 2022). For instance, CT images predicted from MRI scans are useful for radiotherapy planning (Boulanger et al., 2021), PET images generated from CT scans enhance liver lesion segmentation (Ben-Cohen et al., 2019), and generated X-rays can be used to reduce radiation dose in cardiac interventions (Azizmohammadi et al., 2022). Other methods aim to harmonize images from different scanners by converting them to a unified style or developing domain-agnostic models to maintain good performance despite domain shifts – thus facili-

tating subsequent harmonization or even rendering it unnecessary (C. Ma et al., 2019; Yamashita et al., 2021).

Similarly, synthetic corruption and downsampling of real data is useful in applications like image quality transfer (IQT), which estimates high-quality data from low-quality images, in a deterministic or probabilistic fashion. While IQT with paired low- and high-quality images acquired separately (e.g., on different devices) is possible, IQT based on synthetic downsampling is far more common as it only requires the high-quality scans (and also bypasses the need to accurately co-register the image pairs). IQT enables the automatic improvement of images acquired with portable or older MRI scanners with lower field strength (Alexander et al., 2017; Blumberg et al., 2018; Jones et al., 2018; Kim & Alexander, 2021; C.-H. Lin et al., 2018; Tanno et al., 2020). Several works have shown that combining real and synthetic scans endows IQT (both in its deterministic and probabilistic version) with increased generalization ability at test time (Figini et al., 2020; Kim et al., n.d.; H. Lin et al., 2023).

2.3 Methods for synthesizing data in medical image analysis

Model-based techniques, one category of medical image synthesis, rely on hand-crafted pipelines to generate data based on known physical properties of the desired image distributions.The simplest of these approaches uses geometric and intensity augmentations of existing images or label maps (Shorten & Khoshgoftaar, 2019). Geometric augmentation samples a spatial transform that deforms anatomy, while intensity augmentation randomly varies the image intensity profile – for example, by randomly modifying the brightness, contrast, and gamma correction of the images; or by filtering them using random convolutional kernels (Xu et al., 2020) or shallow networks with random weights (Ouyang et al., 2022). More complex techniques leverage biophysical models of the imaging signal to simulate realistic acquisitions or use deformable image registration to induce longitudinal growth or atrophy within an existing scan (Graff, 2016; Jog et al., 2019; Kainz et al., 2019; Karaçali & Davatzikos, 2006; Khanal et al., 2017; Larson & Oguz, 2022; Sengupta et al., 2021; A. D. C. Smith et al., 2003; Teixeira et al., 2018).

Alternatively, recent generative deep-learning methods that learn to replicate spatial patterns in a training distribution are used to produce synthetic images that resemble real-world data. Even when actual data exist, these models can help in scenarios where data sharing is restricted due to privacy concerns, as synthetic images can often be disseminated with fewer limitations (Pinaya et al., 2022). Deep generative model architectures like generative adversarial networks (GANs) (Bowles et al., 2018; Y. Chen et al., 2022), variational autoencoders (VAEs) (Casale et al., 2018; Kingma & Welling, 2013), normalizing

flows (Wilms et al., 2022), and diffusion probabilistic models can achieve high-resolution image synthesis (Chung et al., 2022; Ho et al., 2020; Kazerouni et al., 2023; Rombach et al., 2022). Conditional generative models can also be used to produce or augment images based on a prompt. For instance, given a label map, BrainSpade (Fernandez et al., 2022) generates brain scans conditioned on features of reference scan to control generated image style or the presence of certain pathology. RoentGen (Chambon et al., 2022) synthesizes chest X-rays with features based on an input natural language description.

A related approach to fully generative modeling is semi-supervised learning, where the same label maps for generative modeling are used to facilitate learning from numerous unlabeled examples and a small number of labeled ones. Semi-supervised learning approaches are particularly useful if the generative forward model is complex or otherwise challenging to simulate, leading to a large gap between synthesized and real data. Approaches that employ label maps for semi-supervision include supervision by denoising (SUD) (Young et al., 2023) and denoising diffusion probabilistic model (DDPM) variants (Chan et al., 2023) with applications to whole brain segmentation and cortical parcellation.

We emphasize, however, that the capacity of generative methods to synthesize diverse images is ultimately confined by the scope of data used to train it. As a result, the robustness of downstream models trained on generated outputs is constrained not only by the overall accuracy of the generator, but also by potential domain gaps that exist in real data.

3 Brain image synthesis using domain randomization

To address training domain gaps in learning-based brain imaging methods, our group has applied a model-based approach for constructing entirely synthetic training datasets that feature anatomies, intensity distributions, and artifacts that extend well beyond the realistic range of medical images. This strategy facilitates training models capable of generalizing across a variety of *real* brain images.

3.1 Method

This framework requires a pre-computed set of whole-brain, anatomical label maps $\mathcal{S}$, derived manually or automatically. To generate an image, we sample a label map $s \sim \mathcal{S}$ and apply a domain randomization model following Figure 1. First, we generate a transformed map s_ϕ with varied anatomical morphology by applying to s a random affine and non-linear deformation pair ϕ. Next, we consider a model of tissue contrast inspired by the classical Bayesian segmentation literature (Ashburner & Friston, 2005)

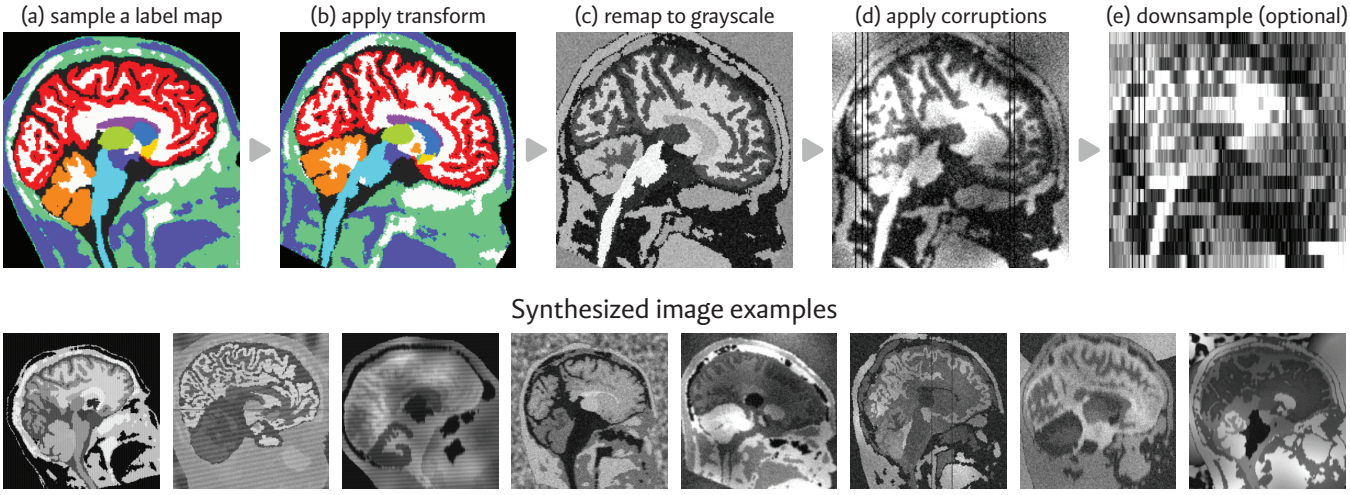

Figure 1: Top: Overview of our image synthesis process. First, we (a) sample a prior-generated whole-brain label map and (b) transform it with random affine and deformable spatial augmentations. Then, we (c) remap the labels to grayscale, drawing from random Gaussian distributions, and (d) simulate artifacts, such as noise, intensity bias, and smoothing. In some applications, we (e) resample the image to a new voxel spacing. Bottom: Series of images generated with random synthesis parameters, each derived from the *same* initial label map.

and generate an intensity distribution for each label in s_ϕ, using uniformly sampled Gaussian parameters. With this random mixture model, we compute a grayscale image x by recoding voxel labels in s_ϕ with values sampled from their corresponding label-specific intensity distribution. Lastly, we augment x with randomly simulated image artifacts such as spatial blurring, added noise, bias field exponentiation, and slice corruption. In some cases, we manipulate the acquisition geometry, for instance by randomly cropping the field-of-view or resampling to a new voxel spacing in a sampled slice direction (axial, coronal, or sagittal).

We follow the standard deep learning setup, but instead of using acquired brain data, we use the synthetic images during training. As usual, we propagate the prediction error using ground truth annotations or an unsupervised loss function. For example, in some applications, components of the deformed label map s_ϕ can be used as targets to train a segmentation model. Alternatively, ϕ can be applied to original acquisitions to compute targets for an image reconstruction model.

3.2 Downstream model robustness

Brain image analysis models trained using this approach are broadly applicable and generalize well across MRI sequences and contrasts, field strengths, and scanner manufacturers – all without using real acquisitions during training or having to retrain networks for different domains. This generalizability enhances

the utility of neuroimaging models as robust, deployable tools that clinicians and researchers alike can readily apply to their own data, without requiring specialized hardware or machine learning expertise to fine-tune inflexible models for particular acquisition specifics.

Additionally, compared to supervised models developed with real images, those optimized with synthetic data and training targets derived from s_ϕ are less vulnerable to imperfections in ground-truth data, due to the inherent alignment between label maps and synthesized images. While not entirely resilient to major segmentation errors, this method mitigates common issues using conventional data, where small discrepancies between target labels and underlying image features often disrupt the learning process. Consequently, it minimizes the need to produce highly precise, manual segmentations for training.

We emphasize that while our approach improves model robustness across acquisition specifics, it may not completely bridge the domain gap for certain intensity and spatial features that are not encompassed within both the synthesis protocol and training cohort. For instance, a model optimized with synthetic data exclusively derived from scans of healthy adults may struggle to adapt to images featuring significant pathology or those from infants. In section 5, we elaborate on our ongoing and future efforts to improve and expand our framework for enhanced model applicability for diverse imaging conditions and populations.

4 Current applications

Over the past few years, our group has successfully applied this synthesis strategy to tackle several neuroimaging problems. This section reviews these deployed applications, illustrated in Figure 2.

4.1 Whole brain segmentation and skull stripping

SynthSeg, the first tool to encapsulate our synthesis-based strategy, is a contrast-agnostic model that segments any brain MR scan into 32 anatomical regions (Billot, Greve, et al., 2020, 2023; Billot, Robinson, et al., 2020). We optimize SynthSeg using synthetic images as input and components of the deformed source label map s_ϕ as the target output (Figure 2a). By randomly downsamping and smoothing synthetic inputs, we train SynthSeg to produce 1 mm isotropic segmentations even for low-resolution images at inference; we note that, while this output resolution is constant, the predictions are expected to be crisper and more accurate for higher-resolution input scans. A subsequent iteration of SynthSeg (Billot, Magdamo, et al., 2023) includes four crucial improvements that enable neuroimaging

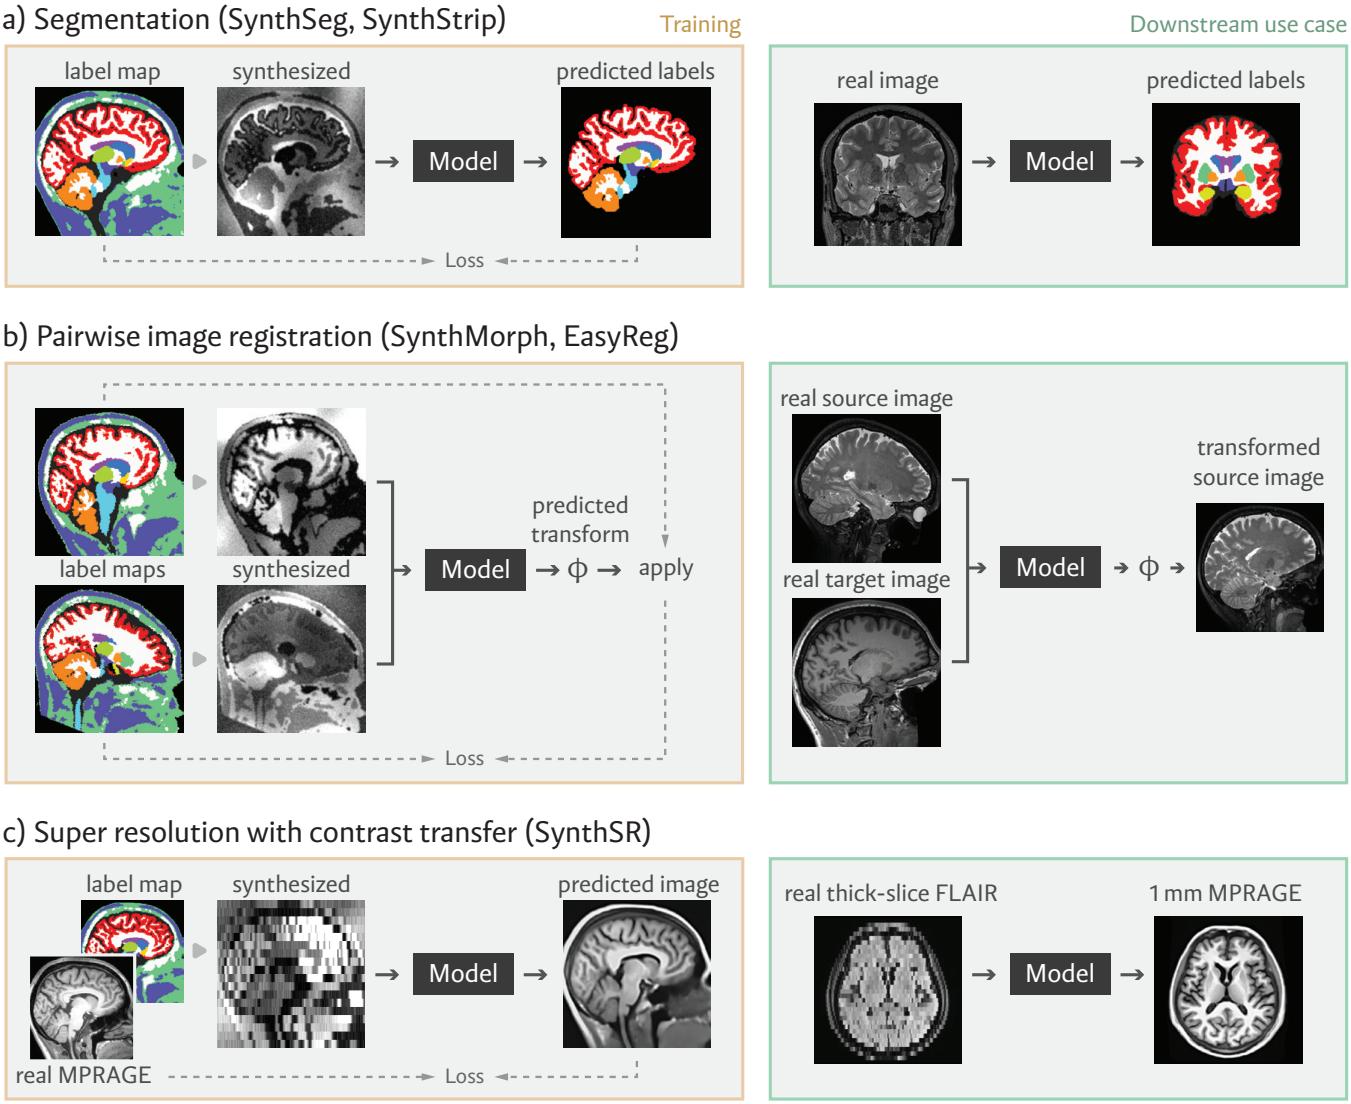

Figure 2: Applications using our synthesis framework. These include robust, contrast-agnostic methods for (a) whole-brain segmentation and skull-stripping, (b) linear and non-linear alignment, and (c) super-resolution and contrast transfer to MPRAGE characteristics. For each tool, orange boxes on the left illustrate the model optimization procedure using synthetic data, and green boxes on the right demonstrate downstream use of the model at inference time.

studies of uncurated clinical MRI data directly downloaded from the picture archiving and communication system (PACS) of a hospital. First, a hierarchical approach greatly enhances the robustness of the method by segmenting four tissue types. We ensure the plausibility of the tissue segmentation with an autoencoder and segment the brain regions conditioned on the tissue types. Second, we subdivide of the cortex into parcels. Third, we estimate the intracranial volume, an important covariate in volumetric studies. Fourth, we include an automatic quality control module that automatically rejects poorly segmented scans.

SynthStrip is another contrast-agnostic segmentation model that we release (Hoopes, Mora, et al., 2022;

Kelley et al., 2024). This tool provides a binary image delineation of brain and non-brain voxels – a task commonly known as skull stripping or brain extraction (Eskildsen et al., 2012; Iglesias et al., 2011; Ségonne et al., 2004; S. M. Smith, 2002). This is a fundamental first step in many neuroimaging pipelines as it removes irrelevant features from the image and facilitates downstream analyses like brain-specific registration. During training, SynthStrip learns to predict a signed distance map to the brain matter boundary, and unlike earlier deep learning-based skull stripping algorithms (Kleesiek et al., 2016; Roy et al., 2017; Salehi et al., 2017), we leverage the synthesis strategy to achieve robust brain extraction across a variety of data types.

4.2 Image registration

Medical image registration estimates a spatial transformation between corresponding anatomies of two images. Classical methods rely on iterative optimization approaches (Ashburner, 2007; Avants et al., 2008; Modat et al., 2010; Rueckert et al., 1999), whereas modern learning-based registration methods directly predict a deformation from two input scans using a neural network. A number of early learning registration approaches optimize a loss comparing the network output to a ground truth warp (Krebs et al., 2017; Rohé et al., 2017; Yang et al., 2017; Young et al., 2022). More recently, unsupervised methods train with image similarity functions similar to those of classical methods (Balakrishnan et al., 2019; De Vos et al., 2019; Hoopes, Hoffmann, et al., 2022; Krebs et al., 2019). While these techniques yield high accuracy in intra-modality registration, they inherit the limitations of classical techniques (lack of robustness) when registering across modalities (Iglesias et al., 2013).

In SynthMorph (Hoffmann, Billot, et al., 2021; Hoffmann et al., 2022, 2023, 2024) and its extension EasyReg (Iglesias, 2023), we train a registration network with synthetic images of random resolution and contrast, circumventing the need for inter-modality similarity measures. Specifically, these methods rely on generating two, rather than one, synthetic images from the same label map, using different random deformations, resolutions, and appearances (Figure 2b). The network can then be trained to predict deformations in a supervised or unsupervised fashion. The former relies on the known ground truth deformation similar to Young et al., 2022, whereas the latter maximizes a measure of structure overlap, such as the Dice score, using the source label maps ϕ. EasyReg extends SynthMorph by estimating a deformation field in both registration directions. The forward deformation is averaged with the inverse of the reverse deformation to guarantee bi-directional symmetry.

We have also explored the possibility of training SynthMorph with label maps drawn from discrete, spatially smooth noise distributions, such as Perlin noise (Hoffmann et al., 2022). A network trained

with such data enables registration of any anatomical region, at the expense of a slight loss of performance with respect to the dedicated brain model when registering brain MRI scans. Even for imaging modalities or body parts where the model may underperform, it is still useful as an unbiased starting point to initialize model training (Dey et al., 2022).

4.3 Super-resolution and contrast transfer

Synthetic images can also be used to solve other voxel-wise regression problems, such as super-resolution and contrast transfer. A super-resolution network (Iglesias et al., 2020, 2022) can be trained using low-resolution synthetic images as input and the corresponding high-resolution synthetic images as the targets. A similar strategy can be used to train networks for MRI contrast transfer (e.g., FLAIR to T1w), using synthetic images produced with our generator. SynthSR (Iglesias et al., 2023) achieves super-resolution and contrast transfer simultaneously (Figure 2c), producing 1 mm isotropic T1w volumes from brain MRI scans of any slice direction, resolution, and contrast. This synthetic T1w is compatible with most existing morphometric tool for brain MRI analysis.

Training SynthSR requires real 1 mm T1w scans matching the label maps. In practice, this is not prohibitive since label maps are typically obtained from real T1w scans on healthy subjects in the first place. These T1w images are deformed by ϕ and used as regression targets by a neural network that seeks to predict them from synthetic input scans of random resolution and contrast. We optimize the network with a primary similarity loss using the sum of absolute image intensity differences between the prediction and the ground truth. When used in isolation, this loss yields predictions that are blurrier than ideal. To address this, we employ a second loss based on the Dice overlap between s_ϕ and the automated segmentation of the predicted images, estimated by a segmentation network. The weights of this second network are frozen during training, thus encouraging SynthSR to generate images that are well segmented and thus well defined at anatomical boundaries.

5 Ongoing and future work

In this section, we discuss ongoing and future efforts to apply our synthetic framework and various aims to enhance the effectiveness of synthesized images.

5.1 Improving synthesized images

Synthesizing image intensity features by recoding anatomical label maps may not fully capture the nuanced characteristics that exist in real acquisitions, as the anatomical detail of generated images are constrained by the granularity of the underlying segmentations. In some cases, this limitation might lead to trained models that do not fully match the within-domain performance of a model trained on a *specific* acquisition type. While one solution is to leverage high-resolution, densely-segmented maps with hundreds of anatomical regions, this is exceedingly difficult to acquire manually or even through automated means. A less prohibitive solution involves integrating real data and alternative synthesis methods alongside our framework to further expand the distribution of tissue characteristics seen by the model during training. This could involve occasionally applying different synthesis techniques (or none at all) or combining multiple synthesized images from various techniques into a single, joint image.

One useful alternative synthesis method capitalizes on MRI signal models to simulate *realistic* images with diverse acquisition characteristics via tissue parameter maps (Jacobs et al., 2023; Varadarajan et al., 2021). Specifically, pulse sequences like DESPOT2 (Deoni et al., 2005), MP2RAGE (Marques et al., 2010), and multi-echo GRE (S. Shin et al., 2023) compute magnetic property maps of the tissue, such as T1 and T2* relaxation time and proton density (PD). These maps can be combined as input to a forward model derived from the Bloch equation (Bloch, 1946; D. Ma et al., 2013) to generate images that closely mirror the appearance of a real acquisition. By randomly varying sequence parameters in the Bloch simulation, we can generate realistic images with a range of tissue contrasts, all from a single acquisition. However, with this technique, images are not computed directly from a label map and therefore require more precise image annotations to train accurate models in supervised scenarios.

Another approach to synthesis involves generative modeling techniques that learn to reconstruct new brain images from input label maps (Fernandez et al., 2022; Pinaya et al., 2022). Looking forward, multi-modal neuroimaging models could further augment our framework with fine-grained control, for instance by using a natural language prompt to specify desired image properties or alterations. While these models tend to interpolate between features seen in real data and do not substantially exceed characteristics of the training domain, they provide a way to generate detailed images from labels used to supplement our framework.

A related avenue of future research is conditional generation of synthetic scans (Chambon et al., 2022; Fernandez et al., 2022; Pinaya et al., 2022), given non-imaging variables such as age and gender. Modeling such variables may not be trivial (e.g., how does the "age" of a brain scan change if a random nonlinear augmentation expands the ventricles?), but would open the door to applications such as age

prediction ("BrainAge"), which is a reliable biomarker for cognitive impairment and several diseases(Cole et al., 2017).

5.2 Modeling imaging artifacts

By simulating acquisition errors and distortions, we can potentially train models that learn to extract and correct these corruptions in real images. In MRI, the acquisition is performed in the frequency domain (k-space), and errors in the frequency domain can have a rather broad effect on the image, such as localized distortions or spurious sinusoidal patterns (Graves & Mitchell, 2013; Jezzard & Balaban, 1995). Modern MRI acquisitions also often undersample k-space through the use of parallel imaging or compressed-sensing (Griswold et al., 2002; Lustig et al., 2008; Pruessmann et al., 1999), which require more advanced reconstruction algorithms that cause specific artifacts (Brau et al., 2008). The use of slice-selective pulses can also lead to spin-history artifacts, caused by imperfect pulse shapes (Malik et al., 2011) or head motion (Friston et al., 1996) – which itself can also disrupt the consistency of k-space and lead to smearing and duplicated edges (Zaitsev et al., 2015). Each of these phenomena can be synthesized in brain images, either using rough approximations or by simulated physical processes with supplemental data.

Synthetic images also have use for bias field estimation and correction of both B1- (receive) and B1+ (transmit) inhomogeneities. The receive field is an intensity scaling across space and easily simulated in its general appearance as a spatial distribution of high sensitivity regions near coils with a rapid fall-off depending on the size of the coil (so higher sensitivity scaling also results in more rapid fall-off from smaller coils). While this type of profile is simple to approximate in image-based augmentation, doing so yields an image with a bias field that is the composition of the true bias field and the simulated one.

In contrast, one can include the receive bias field directly into the image synthesis procedure so it is the only one present in the resulting data. On the transmit side, the steady state Bloch equations (Bloch, 1946) can be integrated into the synthesis, as long as parameter maps are available. In that case, the transmit bias field is applied to a randomly generated flip angle to yield a "flip angle field" that modulates the contrast of the image in a tissue-dependent fashion through the Bloch equations. A receive bias field can then be applied after the initial image formation from the magnetic properties.

While modeling the precise form of the B1+/- fields is difficult, simulating a wider range of variation than we expect to see in practice may be sufficient to learn to disentangle and separately predict the fields. Alternatively, one could also use an adversarial approach to simulate realistic bias fields that maximize the performance of a downstream tasks, e.g., image segmentation (C. Chen et al., 2020).

5.3 Targeted anatomical analysis

Brain synthesis methods promise to improve existing anatomy-specific analysis approaches in various contexts.

5.3.1 Surface reconstruction

Cortical reconstruction techniques seek to localize and model the tissue boundaries of the white matter and pial surfaces, to facilitate accurate analysis of the morphological, functional, and connective organization of the cortex. Current tools like FreeSurfer rely on time-consuming, per-image optimization methods (Dale et al., 1999; Fischl et al., 1999), but recent learning-based approaches offer solutions to efficient and accurate reconstruction (Bongratz et al., 2022; Cruz et al., 2021; Q. Ma et al., 2021, 2022). To facilitate surface fitting across a range of image contrasts, we are augmenting our cortical reconstruction tools (Gopinath et al., 2023, 2024; Hoopes, Iglesias, et al., 2022) by leveraging a mix of real and synthetic images during training. Our straightforward label-to-image synthesis does not reproduce the intensity features necessary for proper reconstruction of highly-folded cortical convolutions. To address this, we are currently leveraging more realistic synthesis methods based on parameter maps, as described in section 5.1.

5.3.2 Brain vasculature

The vascular anatomy is of great interest both for research and medical applications (stroke, atherosclerosis, aneurysm). Arteries are typically imaged using variants of MR angiography (contrast-based, time-of-flight), while veins are detected using susceptibility imaging (susceptibility-weighted imaging, quantitative T2*, quantitative susceptibility mapping). Finer levels of neurovasculature can be visualized in high-resolution *post mortem* MRI or using microscopy techniques.

Vascular modeling therefore requires the acquisition, segmentation, and analysis of multimodal, multiscale, cross-sectional data. Vessel segmentation has historically relied on Hessian filters that leverage the specifically anisotropic nature of image curvature around vessels (Frangi et al., 1998; Lorenz et al., 1997; Sato et al., 1998). However, these filters are local, sensitive to noise, and require tedious parameter tuning. Deep learning has shown promise, but vessel contrast varies greatly across scales and modalities, thereby requiring relabeling (which is particularly tedious and thus expensive) and retraining for each new case.

In this context, the synthesis of heterogeneous training data for vascular analysis is extremely appealing.

There is a large body of literature concerned with the generation of realistic synthetic vascular networks, to accurately model the dynamics of blood flow and blood oxygenation. An early computation method is "constrained constructive optimization" (CCO), originally developed by Neumann et al. (1995) and later refined by their group and others (Karch et al., 1999, 2000; Schneider et al., 2012; Szczerba & Székely, 2005). Its use for training and validating segmentation algorithms was first proposed in (Hamarneh & Jassi, 2010) and was recently applied to the problem of retinal OCT segmentation (Menten et al., 2022). We believe that combining domain randomization techniques with vascular simulators has great potential in enabling vascular segmentation in a wide array of neuroimaging applications.

5.3.3 Change detection and longitudinal analysis

Longitudinal within-subject image analysis approaches have the potential to increase the sensitivity and specificity of population studies. These methods can improve the efficiency of trials by requiring fewer subjects and providing surrogate endpoints to assess therapeutic efficacy. Current analysis tools effectively detect longitudinal changes in well-curated research data, such as ADNI. In these datasets, scan protocols are harmonized across acquisition sites to minimize differential distortions, and residual distortions, such as gradient nonlinearities, are corrected before data release.

Unfortunately, these tools fail in the presence of acquisition differences between scans. Such differences are ubiquitous in imaging collections, and especially prevalent in clinical imaging, where scheduling a subject on the same scanner and scan protocol as a previous session is difficult or impossible. The tools perform poorly in many challenging scenarios where they are faced with complex noise, anatomical atrophy, and varying MRI contrast and distortion across serial scans. It is therefore important to develop tools that can ignore large-scale technology-induced differences in the longitudinal scans, while finding subtle anatomical changes that indicate early disease processes such as atrophy in Alzheimer's Disease (AD).

To overcome these barriers, the synthesis strategy presents an opportunity to detect potentially subtle neuroanatomical change, in the presence of large image differences coming from (uninteresting) aspects of the acquisition process, including field strength, receive coil, sequence parameters, gradient nonlinearities, and B0-distortions/read-out directions. Such artifacts make it challenging to obtain ground truth in real datasets. We believe that the synthesis strategy can be expanded to the longitudinal case, where both the target neuroanatomical changes and the nuisance effects can be modelled during image generation, and hence separately predicted during network training (Bernal et al., 2021; Karaçali & Davatzikos, 2006; Khanal et al., 2017; Larson & Oguz, 2022; Rusak et al., 2022). This strategy promises to enable longitudinal analysis and change detection from large-scale, potentially clinical-quality, image collections

that was not previously possible.

5.3.4 Brain development and myelination

The morphology of the brain quickly changes during childhood, followed by a long period of very slow changes during adulthood (Bethlehem et al., 2021). However, medical image processing pipelines are set to work on normative adult brains and overlook challenges present in early and late life. During childhood, maturation of different brain regions takes place at different periods and rates, yielding large differences in structure size, MRI signal, and image contrast between development stages and compared to adult brains (Dubois et al., 2021).

One example of such change is myelination, the complex process of insulating neuronal axons, which enables rapid communication across the brain. Myelination varies across developmental stages (Monje, 2018), brain regions (Nieuwenhuys & Broere, 2017), and individuals (Van Essen & Glasser, 2014), and is also affected by various neurological disorders (Fields, 2008). Brain MRI analysis of myelinating brains in prenatal and infant MRI is particularly challenging, for two reasons: the progressive flip in image contrast between gray and white matter, and the multi-modal distribution of white matter intensities due to varying levels of myelination. Synthetic training data offer an attractive solution to this problems (Shang et al., 2022): the contrast flip is modeled by the random distribution of Gaussian means, whereas the spatially varying myelination of the white matter can be modeled by subdividing this tissue type into sub-labels with different myelin levels – which can be done automatically by applying clustering algorithms to the intensities of the training images.

5.3.5 Ex vivo MRI

The analysis of *Ex vivo* imaging data (Augustinack et al., 2005; Edlow et al., 2019) presents unique challenges due to variability in anatomical structures and appearances from changes in the packing container and liquid. For example, the image may contain one hemisphere or both hemispheres, either with or without a brainstem or cerebellum. Cerebral vessels can contain residual iron content in some locations and not others, varying the contrast properties of the vasculature. Proton-free fluids such as fomblin can be used for packing so that the exterior is dark, or water-based fluids can be used, resulting in a bright exterior in many MRI contrasts. These effects are difficult to model in standard image-based augmentation. Nevertheless, in our synthesis framework, neural networks can be trained to be robust to the absence of structures by probabilistically removing or retaining during training.

5.4 Synthesis of brain abnormalities

The use of inpainting methods to remove pathology from disease-case data, such as white matter lesions or brain tumors (Iglesias et al., 2023; Prastawa et al., 2009; H. Zhang et al., 2020), or to inject synthetic lesions into scans of healthy subjects (da S Senra Filho et al., 2019). The former is advantageous for generating synthetic, "healthy" data that can be analyzed with image processing methods that may fail in the presence of severe abnormalities; the latter can be used to generate large amounts of data with known ground truth.

In the context of our synthesis framework, we have shown that our random Gaussian model combined with existing segmentations is sufficient to capture the appearance of white matter hyper- or hypo-intensities due to multiple sclerosis lesions (Billot et al., 2021). Integrating other types of abnormality with improved models of appearance has the potential to enable the automated analysis of large amounts of heterogeneous data from PACS around the world, with minimal or no curation. While modeling the appearance of common pathologies like stroke or tumors is certainly more difficult than for white matter hyperintensities, the availability of large public datasets with manual delineations (Liew et al., 2017; Menze et al., 2015) should facilitate this endeavor. For more uncommon pathologies, the absence of manual labels is the main obstacle to training machine learning models. An appealing solution consists of generating such labels from scratch, using random shapes Hoffmann et al., 2022, geometric rules (Dey et al., 2023), or more advanced constructive synthesis (Georg et al., 2010; Menten et al., 2022; Schneider et al., 2012), with the hope that the synthesized data make the network robust against the presence of pathologies that are not present in the training dataset.

5.5 Applications in neuropathology

Synthesizing training data has potential applications for neuroimaging to neuropathology correlation. These studies aim to correlate gold standard histopathological diagnoses and microscopic measurements with macroscopic biomarkers and morphometry (Charidimou et al., 2020; Edlow et al., 2018; Nolan et al., 2021; Webster et al., 2021). In the absence of timely ante mortem MRI or cadaveric imaging, such studies rely on *ex vivo* MRI scans for which training data are scarce. As explained above, our synthesis framework allows the construction of *ex vivo* analysis tools that benefit from the wealth of data available for *in vivo* MRI.

In a similar vein, such training data can be applied to new volumetric imaging modalities. For example, one proposed solution to the difficulty of image acquisition for neuroimaging to neuropathology correlation is the construction of 3D volumes from abundant 2D dissection photography (Gazula et al., 2023; Tregidgo

et al., 2020). Modifying resampling and bias augmentations to account for 2D slices during training allows the generalization of existing domain randomization tools to these images, without the need for a large training set from an entirely new modality.

5.6 Universal models within and beyond neuroimaging

In most existing frameworks we discussed, synthesis is used to build robust models tackling a *single* task, such as registration or segmentation of specific regions of interest. More recently, more *general* neuroimaging models also employ these synthesis strategies but aim to tackle multiple tasks at once. For example, Neuralizer is an in-context learning framework that can solve new neuroimaging tasks at inference (Czolbe & Dalca, 2023). As input, it takes an image to be processed along with pairs of input-output examples that illustrate the task to be executed. Neuralizer employs synthetic images derived with our framework to ensure robustness to new tasks as well as new input and output image modalities and qualities. Additionally, synthesis can be used to train robust, multi-task models that learn to extract broadly useful neuroimaging features, functioning as initial training checkpoints for rapid fine-tuning of single-task objectives (Chua & Dalca, 2023; P. Liu et al., 2023) Finally, these approaches can be expanded beyond brain imaging, for example, by using label maps of random shapes to train generalized medical imaging tools. Existing methods achieve anatomy-invariant solutions for registration (Hoffmann et al., 2022, 2024), in-context segmentation (Butoi et al., 2023), interactive segmentation (Wong et al., 2023), and star-convex instance segmentation (Dey et al., 2023).

6 The curse of dimensionality and modeling limitations

Brain MRI scans that are acquired in clinical routine imaging are not only variable in terms of intensity contrast and resolution properties – which many of our synthesis-based tools successfully address – but also along several other dimensions: number and type of contrasts that are acquired, administration of a contrast agent, brain disorders not seen during training (as explained in section 5.4), or the number of timepoints (and time interval between scans) in longitudinal imaging.

While data synthesis can conceivably cover each of these dimensions of variation when taken individually, training models that address all of them *simultaneously* incur the "curse of dimensionality", the amount of data that needs to be synthesized grows *exponentially* with the number of dimensions of variation that the models should be robust against.

Furthermore, differences in acquisition and reconstruction strategies (e.g., compressed sensing or parallel imaging) can yield correlations across spatial locations and spatially varying noise distributions (Breuer et al., 2009; Griswold et al., 2002; Lustig et al., 2008; Pruessmann et al., 1999), which are currently not captured by our simplified model. In a similar manner, advanced machine learning MRI reconstruction methods (e.g., AUTOMAP, Zhu et al. 2018) generate highly structured and signal-dependent noise that cannot be modeled analytically. These differences in noise distribution may disturb neural networks, especially when the signal-to-noise ratio is low, as in low-field MRI. More importantly, in the context of quantitative MRI (relaxometry, diffusion, perfusion), quantitative estimates can become heavily biased if the correct distribution is not taken into account, either during data sampling or in the loss calculation (Sijbers et al., 1999; Varadarajan & Haldar, 2015).

Tackling these problems will likely require training with synthetic data generated with advanced simulations (as discussed in Section 5.1 above), which will most probably be too computationally expensive to run on the fly on a common server. Circumventing this limitation may require supercomputing power, offline generation of a large (but finite) number of samples, or the development of computationally efficient approximations that yield comparable performance.

7 Discussion and conclusion

In this article, we have reviewed our recently proposed paradigm to train neural networks with synthetic data, along with several applications that benefit (or have the potential to benefit) from the framework.

One area of potential impact is clinical neurology and translational neuroscience (Goetz, 2007; Matthews et al., 2006; Simon et al., 2009), via analysis of clinical *in vivo* MRI scans C. Zhao et al., 2020. For the broad spectrum of neurological disorders associated with focal brain lesions, such as traumatic brain injury, ischemic stroke, intracerebral hemorrhage, or encephalitis, our understanding of the anatomic correlates of cognitive, physical, and behavioral symptoms has been hindered by an inability to process anisotropic images or lesioned brains in neuroimaging software platforms (Iglesias et al., 2023). Lesions often compromise the accuracy of cortical surfaces that are essential to generate volumetric measurements, leading to the exclusion of lesioned brains from clinico-radiologic correlation studies (Nuesch et al., 2009). Fundamental problems in neuroscience with profound implications for clinical neurology, such as understanding neural correlates of consciousness, remain unanswered, partly because precise MRI mapping of lesions is often infeasible in patients with disorders of consciousness (Edlow et al., 2024). By making it possible to generate cortical surfaces in images of lesioned brains acquired with arbitrary resolution and contrast, our techniques can create new opportunities to map lesions onto canonical

brain atlases in 1 mm stereotactic space – thereby identifying the precise anatomic correlates of a broad spectrum of neuropsychological symptoms (Johnson et al., 2012; Montoya et al., 2006; Politis, 2014).

Our tools also have the potential to create opportunities to perform retrospective studies that mine massive clinical datasets whose potential has not been tapped. Clinical MRI scans that were previously only amenable to qualitative analysis can now be tested for associations between cortical and subcortical volumetric measures and clinical syndromes. Moreover, a large number of patients with neurodegenerative disorders worldwide have had serial clinical MRI scans during decades of care (Rinck, 2019). The integration of our tools with longitudinal analysis pipelines may generate insights into patterns of brain atrophy that have not been possible in prospective studies. In short, our tools can create opportunities to retrospectively study clinical MRI datasets, regardless of contrast, spatial resolution, or lesion burden, with sample sizes that are orders of magnitude larger than those of current prospective studies.

Finally, these strategies can impact domains where obtaining accurately labeled data is infeasible, such as longitudinal studies of cortical atrophy, or vessel segmentation (Chollet et al., 2024; Larson & Oguz, 2022; Rusak et al., 2022). These data hold promise in enabling supervised training of neural networks that achieve superhuman accuracy in these tasks (Azizi et al., 2023; Nikolenko, 2021). As the quality of simulations, the availability of data, and computational capabilities grow, learning-based methods relying on synthetic imaging data promise to play an important role in robust neuroimage processing tools.

Author Contributions

Conceptualization: KG, AH, AVD, JEI. Writing - Original Draft: All authors. Writing - Review & Editing: KG, AH, AVD, JEI. Supervision: AVD, JEI.

Funding

Support for this research was provided in part by the National Institute for Biomedical Imaging and Bioengineering (P41 EB030006, R01 EB031114, R01 EB033773), the National Institute on Aging (R21 AG082082, R01 AG064027, R01 AG070988, RF1 AG080371, K99 AG081493, P30 AG066509, P30 AG066567, U19 AG060909), the National Institute of Mental Health (U01 MH117023, UM1 MH130981, RF1 MH121885, RF1 MH123195), the National Institute of Child Health and Human Development (R00 HD101553), the National Institute for Neurological Disorders and Stroke (R01 NS083534, U24 NS100591, R01 NS105820, U24 NS135561, U01 NS137484, U01 NS132181, UM1 NS132358),

the Brain Initiative Cell Census Network, the Brain Initiative Brain Connects Consortium, the Massachusetts Life Sciences Center, the Universitat de Girona (IFUdG2023), the Medical Research Council (MR/W031566/1), the UCLH Biomedical Research Centre, the UCL Centre for Doctoral Training in Intelligent Integrated Imaging in Healthcare (EP/S021930/1), the Engineering and Physical Sciences Research Council (EP/Y028856/1), the Wellcome Trust (Award 221915), the Lundbeck Foundation (R360–2021–39), the Nancy and Buster Alvord Endowment, the Department of Defense (W81XWH-21-S-TBIPH2), and the Alzheimer's Association (24AARG-NTF-1187394).

Declaration of Competing Interests

Multiple authors maintain consulting relationships with DeepHealth (AVD, BF), SidestepAI (AVD), Neuro42, Inc (MH), Astrocyte Pharma (WTK), Acasti Pharma (WTK), and Ji Xing Pharma (WTK). WTK is partially funded by Hyperfine, Inc. These organizations had no involvement in the preparation or content of this retrospective work.

Data and Code Availability

There are no data or code associated with this manuscript.

References

Agn, M., af Rosenschöld, P. M., Puonti, O., Lundemann, M. J., Mancini, L., Papadaki, A., Thust, S., Ashburner, J., Law, I., & Van Leemput, K. (2019). A modality-adaptive method for segmenting brain tumors and organs-at-risk in radiation therapy planning. *Medical image analysis*, *54*, 220–237.

Ahmad, B., Sun, J., You, Q., Palade, V., & Mao, Z. (2022). Brain tumor classification using a combination of variational autoencoders and generative adversarial networks. *Biomedicines*, *10*(2), 223.

Al Khalil, Y., Amirrajab, S., Lorenz, C., Weese, J., & Breeuwer, M. (2020). Heterogeneous virtual population of simulated cmr images for improving the generalization of cardiac segmentation algorithms. *International workshop on simulation and synthesis in medical imaging*, 68–79.

Alexander, D., Zikic, D., Ghosh, A., Tanno, R., Wottschel, V., Zhang, J., Kaden, E., Dyrby, T., Sotiropoulos, S., Zhang, H., & Criminisi, A. (2017). Image quality transfer and applications in diffusion mri. *NeuroImage*, *152*.

Alfaro-Almagro, F., Jenkinson, M., Bangerter, N. K., Andersson, J. L., Griffanti, L., Douaud, G., Sotiropoulos, S. N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., Vidaurre, D., Webster, M., McCarthy, P., Rorden, C., Daducci, A., Alexander, D. C., Zhang, H., Dragonu, I., Matthews, P. M., ... Smith, S. M. (2018). Image processing and quality control for the first 10,000 brain imaging datasets from uk biobank. *NeuroImage*, *166*, 400–424. https://doi.org/https://doi.org/10.1016/j.neuroimage.2017.10.034

Alhaija, H. A., Mustikovela, S. K., Mescheder, L. M., Geiger, A., & Rother, C. (2017). Augmented reality meets computer vision: Efficient data generation for urban driving scenes. *International Journal of Computer Vision*, *126*, 961–972.

Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M. A., Al-Amidie, M., & Farhan, L. (2021). Review of deep learning: Concepts, cnn architectures, challenges, applications, future directions. *Journal of big Data*, *8*, 1–74.

Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *NeuroImage*, *38*(1), 95–113.

Ashburner, J., & Friston, K. J. (2005). Unified segmentation. *Neuroimage*, *26*(3), 839–851.

Augustinack, J. C., van der Kouwe, A., Blackwell, M. L., Salat, D. H., Wiggins, C. J., Frosch, M. P., Wiggins, G. C., Potthast, A., Wald, L. L., & Fischl, B. R. (2005). Detection of entorhinal layer ii using 7tesla [corrected] magnetic resonance imaging. *Annals of neurology*, *57 4*, 489–94.

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.

Azizi, S., Kornblith, S., Saharia, C., Norouzi, M., & Fleet, D. J. (2023). Synthetic data from diffusion models improves imagenet classification. *arXiv preprint arXiv:2304.08466*.

Azizmohammadi, F., Navarro Castellanos, I., Miró, J., Segars, P., Samei, E., & Duong, L. (2022). Generative learning approach for radiation dose reduction in x-ray guided cardiac interventions. *Medical physics*, *49*(6), 4071–4081.

Badano, A., Graff, C. G., Badal, A., Sharma, D., Zeng, R., Samuelson, F. W., Glick, S. J., & Myers, K. J. (2018). Evaluation of digital breast tomosynthesis as replacement of full-field digital mammography using an in silico imaging trial. *JAMA network open*, *1*(7), e185474–e185474.

Bae, G., de La Gorce, M., Baltruvsaitis, T., Hewitt, C., Chen, D., Valentin, J., Cipolla, R., & Shen, J. (2023). Digiface-1m: 1 million digital face images for face recognition. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3526–3535.

Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J., & Dalca, A. V. (2019). VoxelMorph: A learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, *38*(8), 1788–1800.

Ben-Cohen, A., Klang, E., Raskin, S. P., Soffer, S., Ben-Haim, S., Konen, E., Amitai, M. M., & Greenspan, H. (2019). Cross-modality synthesis from ct to pet using fcn and gan networks for improved automated lesion detection. *Engineering Applications of Artificial Intelligence*, *78*, 186–194.

Bernal, J., Valverde, S., Kushibar, K., Cabezas, M., Oliver, A., & Lladó, X. (2021). Generating Longitudinal Atrophy Evaluation Datasets on Brain Magnetic Resonance Images Using Convolutional Neural Networks and Segmentation Priors. *Neuroinformatics*, *19*(3), 477–492.

Bethlehem, R. A. I., Seidlitz, J., White, S. R., Vogel, J. W., Anderson, K. M., Adamson, C. L., Adler, S., Alexopoulos, G. S., Anagnostou, E., Areces-Gonzalez, A., Astle, D. E., Auyeung, B., Ayub, M., Ball, G., Baron-Cohen, S., Beare, R., Bedford, S. A., Benegal, V., Beyer, F., ... Alexander-Bloch, A. F. (2021). Brain charts for the human lifespan. *Nature*, *604*, 525–533.

Billot, B., Cerri, S., Van Leemput, K., Dalca, A. V., & Iglesias, J. E. (2021). Joint segmentation of multiple sclerosis lesions and brain anatomy in mri scans of any contrast and resolution with cnns. *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 1971–1974.

Billot, B., Greve, D. N., Leemput, K. V., Fischl, B. R., Iglesias, J. E., & Dalca, A. V. (2020). A learning strategy for contrast-agnostic mri segmentation. *International Conference on Medical Imaging with Deep Learning*.

Billot, B., Greve, D. N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A. V., & Iglesias, J. E. (2023). Synthseg: Segmentation of brain mri scans of any contrast and resolution without retraining. *Medical Image Analysis*, *86*, 102789. https://doi.org/https://doi.org/10.1016/j.media.2023.102789

Billot, B., Magdamo, C., Cheng, Y., Arnold, S. E., Das, S., & Iglesias, J. E. (2023). Robust machine learning segmentation for large-scale analysis of heterogeneous clinical brain MRI datasets. *Proceedings of the National Academy of Sciences*, *120*(9), e2216399120.

Billot, B., Robinson, E., Dalca, A. V., & Iglesias, J. E. (2020). Partial volume segmentation of brain mri scans of any resolution and contrast. *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part VII 23*, 177–187.

Bloch, F. (1946). Nuclear induction. *Physical review*, *70*(7-8), 460.

Blumberg, S., Tanno, R., Kokkinos, I., & Alexander, D. (2018). Deeper image quality transfer: Training low-memory neural networks for 3d images. *International Conference on Medical Image Computing and Computer-Assisted Intervention*.

Bongratz, F., Rickmann, A.-M., Pölsterl, S., & Wachinger, C. (2022). Vox2cortex: Fast explicit reconstruction of cortical surfaces from 3D MRI scans with geometric deep neural networks. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20773–20783.

Boulanger, M., Nunes, J.-C., Chourak, H., Largent, A., Tahri, S., Acosta, O., De Crevoisier, R., Lafond, C., & Barateau, A. (2021). Deep learning methods to generate synthetic ct from mri in radiotherapy: A literature review. *Physica Medica*, *89*, 265–281.

Bowles, C., Chen, L., Guerrero, R., Bentley, P., Gunn, R. N., Hammers, A., Dickie, D. A., Hernández, M. V., Wardlaw, J. M., & Rueckert, D. (2018). Gan augmentation: Augmenting training data using generative adversarial networks. *ArXiv*, *abs/1810.10863*.

Brau, A. C., Beatty, P. J., Skare, S., & Bammer, R. (2008). Comparison of reconstruction accuracy and efficiency among autocalibrating data-driven parallel imaging methods. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, *59*(2), 382–395.

Breuer, F. A., Kannengiesser, S. A., Blaimer, M., Seiberlich, N., Jakob, P. M., & Griswold, M. A. (2009). General formulation for quantitative g-factor calculation in grappa reconstructions. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, *62*(3), 739–746.

Brumer, I., Bauer, D. F., Schad, L. R., & Zöllner, F. G. (2022). Synthetic arterial spin labeling mri of the kidneys for evaluation of data processing pipeline. *Diagnostics*, *12*(8), 1854.

Butoi, V. I., Ortiz, J. J. G., Ma, T., Sabuncu, M. R., Guttag, J., & Dalca, A. V. (2023). Universeg: Universal medical image segmentation. *Proceedings of the IEEE/CVF International Conference on Computer Vision*.

Casale, F. P., Dalca, A., Saglietti, L., Listgarten, J., & Fusi, N. (2018). Gaussian Process Prior Variational Autoencoders. *Advances in Neural Information Processing Systems*, *31*.

Cha, K. H., Petrick, N., Pezeshk, A., Graff, C. G., Sharma, D., Badal, A., & Sahiner, B. (2020). Evaluation of data augmentation via synthetic images for improved breast mass detection on mammograms using deep learning. *Journal of Medical Imaging*, *7*(1), 012703–012703.

Chambon, P., Bluethgen, C., Delbrouck, J.-B., Van der Sluijs, R., Połacin, M., Chaves, J. M. Z., Abraham, T. M., Purohit, S., Langlotz, C. P., & Chaudhari, A. (2022). Roentgen: Vision-language foundation model for chest x-ray generation. *arXiv preprint arXiv:2211.12737*.

Chan, M., Young, S., & Metzler, C. (2023). SUD$ˆ2$: Supervision by denoising diffusion models for image reconstruction. *NeurIPS 2023 Workshop on Deep Learning and Inverse Problems*.

Charidimou, A., Perosa, V., Frosch, M. P., Scherlek, A. A., Greenberg, S. M., & van Veluw, S. J. (2020). Neuropathological correlates of cortical superficial siderosis in cerebral amyloid angiopathy. *Brain*, *143*(11), 3343–3351.

Chen, C., Qin, C., Qiu, H., Ouyang, C., Wang, S., Chen, L., Tarroni, G., Bai, W., & Rueckert, D. (2020). Realistic adversarial data augmentation for mr image segmentation. *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I 23*, 667–677.

Chen, Y., Yang, X.-H., Wei, Z., Heidari, A. A., Zheng, N., Li, Z., Chen, H., Hu, H., Zhou, Q., & Guan, Q. (2022). Generative adversarial networks in medical image augmentation: A review. *Computers in Biology and Medicine*, *144*, 105382.

Chollet, E., Balbastre, Y., Magnain, C., Fischl, B., & Wang, H. (2024). A label-free and data-free training strategy for vasculature segmentation in serial sectioning oct data. *arXiv preprint arXiv:2405.13757*.

Chua, Y. Z. R., & Dalca, A. V. (2023). Contrast invariant feature representations for segmentation and registration of medical images. *Medical Imaging with Deep Learning, short paper track*.

Chung, H., Lee, E. S., & Ye, J. C. (2022). Mr image denoising and super-resolution using regularized reverse diffusion. *IEEE Transactions on Medical Imaging*, *42*(4), 922–934.

Cole, J. H., Poudel, R. P., Tsagkrasoulis, D., Caan, M. W., Steves, C., Spector, T. D., & Montana, G. (2017). Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *NeuroImage*, *163*, 115–124.

Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research*, *29*(3), 162–173.

Cruz, R. S., Lebrat, L., Bourgeat, P., Fookes, C., Fripp, J., & Salvado, O. (2021). DeepCSR: A 3D deep learning approach for cortical surface reconstruction. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 806–815.

Czolbe, S., & Dalca, A. V. (2023). Neuralizer: General neuroimage analysis without re-training. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6217–6230.

da S Senra Filho, A. C., Simozo, F. H., dos Santos, A. C., & Junior, L. O. M. (2019). Multiple sclerosis multimodal lesion simulation tool (MS-MIST). *Biomedical Physics and Engineering Express*, *5*(3), 035003.

Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage*, *9*(2), 179–194.

Das, A., Xian, Y., He, Y., Akata, Z., & Schiele, B. (2023). Urban scene semantic segmentation with low-cost coarse annotation. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5978–5987.

de Dumast, P., Kebiri, H., Payette, K., Jakab, A., Lajous, H., & Cuadra, M. B. (2022). Synthetic magnetic resonance images for domain adaptation: Application to fetal brain tissue segmentation. *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, 1–5.

De Vos, B. D., Berendsen, F. F., Viergever, M. A., Sokooti, H., Staring, M., & Ivsgum, I. (2019). A deep learning framework for unsupervised affine and deformable image registration. *Medical image analysis*, *52*, 128–143.

Deoni, S. C., Peters, T. M., & Rutt, B. K. (2005). High-resolution T1 and T2 mapping of the brain in a clinically acceptable time with DESPOT1 and DESPOT2. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, *53*(1), 237–241.

Dey, N., Abulnaga, S. M., Billot, B., Turk, E. A., Grant, P. E., Dalca, A. V., & Golland, P. (2023). Anystar: Domain randomized universal star-convex 3d instance segmentation. *arXiv preprint arXiv:2307.07044*.

Dey, N., Schlemper, J., Salehi, S. S. M., Zhou, B., Gerig, G., & Sofka, M. (2022). Contrareg: Contrastive learning of multi-modality unsupervised deformable image registration. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 66–77.

Dima, D., Modabbernia, A., Papachristou, E., Doucet, G. E., Agartz, I., Aghajani, M., Akudjedu, T. N., Albajes-Eizagirre, A., Alnæs, D., Alpert, K. I., Andersson, M., Andreasen, N. C., Andreassen, O. A., Asherson, P., Banaschewski, T., Bargallo, N., Baumeister, S., Baur-Streubel, R., Bertolino, A., ... (KaSP), K. S. P. (2022). Subcortical volumes across the lifespan: Data from 18,605 healthy individuals aged 3–90 years. *Human Brain Mapping*, *43*(1), 452–469. https://doi.org/https://doi.org/10.1002/hbm.25320

Dorjsembe, Z., Pao, H.-K., Odonchimed, S., & Xiao, F. (2024). Conditional diffusion models for semantic 3d brain mri synthesis. *IEEE Journal of Biomedical and Health Informatics*.

Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., & Koltun, V. (2017). Carla: An open urban driving simulator. *Conference on robot learning*, 1–16.

Dubois, J., Alison, M., Counsell, S. J., Hertz-Pannier, L., Hüppi, P. S., & Benders, M. J. (2021). Mri of the neonatal brain: A review of methodological challenges and neuroscientific advances. *Journal of Magnetic Resonance Imaging*, *53*(5), 1318–1343.

Ebner, M., Wang, G., Li, W., Aertsen, M., Patel, P. A., Aughwane, R., Melbourne, A., Doel, T., Dymarkowski, S., De Coppi, P., David, A. L., Deprest, J., Ourselin, S., & Vercauteren, T. (2020). An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain mri. *NeuroImage*, *206*, 116324. https://doi.org/https://doi.org/10.1016/j.neuroimage.2019.116324

Edlow, B. L., Keene, C. D., Perl, D. P., Iacono, D., Folkerth, R. D., Folkerth, R. D., Stewart, W., Donald, C. L. M., Augustinack, J. C., Diaz-Arrastia, R., Estrada, C., Flannery, E., Gordon, W. A., Grabowski, T. J., Hansen, K. E., Hoffman, J. M., Kroenke, C. D., Larson, E. B., Lee, P., ... Dams-O'Connor, K. (2018). Multimodal characterization of the late effects of traumatic brain

injury: A methodological overview of the late effects of traumatic brain injury project. *Journal of neurotrauma*, *35 14*, 1604–1619.

Edlow, B. L., Mareyam, A., Horn, A., Polimeni, J., Witzel, T., Tisdall, M. D., Augustinack, J. C., Stockmann, J. P., Diamond, B. R., Stevens, A., Tirrell, L. S., Folkerth, R. D., Wald, L. L., Fischl, B. R., & van der Kouwe, A. J. W. (2019). 7 tesla mri of the ex vivo human brain at 100 micron resolution. *Scientific Data*, *6*.

Edlow, B. L., Olchanyi, M., Freeman, H. J., Li, J., Maffei, C., Snider, S. B., Zöllei, L., Iglesias, J. E., Augustinack, J., Bodien, Y. G., Haynes, R. L., Greve, D. N., Diamond, B. R., Stevens, A., Giacino, J. T., Destrieux, C., van der Kouwe, A. J. W., Brown, E. N., Folkerth, R. D., ... Kinney, H. C. (2024). Multimodal mri reveals brainstem connections that sustain wakefulness in human consciousness. *Science Translational Medicine*, *16*.

Eskildsen, S. F., Coupé, P., Fonov, V. S., Manjón, J. V., Leung, K. K., Guizard, N., Wassef, S. N., Østergaard, L. R., & Collins, D. L. (2012). Beast: Brain extraction based on nonlocal segmentation technique. *NeuroImage*, *59*, 2362–2373.

Fernandez, V., Pinaya, W. H. L., Borges, P., Tudosiu, P.-D., Graham, M. S., Vercauteren, T., & Cardoso, M. J. (2022). Can segmentation models be trained with fully synthetically generated data? *International Workshop on Simulation and Synthesis in Medical Imaging*, 79–90.

Fields, R. D. (2008). White matter in learning, cognition and psychiatric disorders. *Trends in neurosciences*, *31*(7), 361–370.

Figini, M., Lin, H., Ogbole, G. I., D'Arco, F., Blumberg, S. B., Carmichael, D. W., Tanno, R., Kaden, E., Brown, B. J., Lagunju, I., Cross, H. J., Fernández-Reyes, D., & Alexander, D. C. (2020). Image quality transfer enhances contrast and resolution of low-field brain mri in african paediatric epilepsy patients. *ArXiv*, *abs/2003.07216*.

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

Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., & Dale, A. M. (2002). Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. *Neuron*, *33*(3), 341–355. https://doi.org/https://doi.org/10.1016/S0896-6273(02)00569-X

Fischl, B., Sereno, M. I., & Dale, A. M. (1999). Cortical surface-based analysis: II: Inflation, flattening, and a surface-based coordinate system. *Neuroimage*, *9*(2), 195–207.

Frangi, A. F., Niessen, W. J., Vincken, K. L., & Viergever, M. A. (1998). Multiscale vessel enhancement filtering. *Medical Image Computing and Computer-Assisted Intervention—MICCAI'98: First International Conference Cambridge, MA, USA, October 11–13, 1998 Proceedings 1*, 130–137.

Frangou, S., Modabbernia, A., Williams, S. C. R., Papachristou, E., Doucet, G. E., Agartz, I., Aghajani, M., Akudjedu, T. N., Albajes-Eizagirre, A., Alnæs, D., Alpert, K. I., Andersson, M., Andreasen,

N. C., Andreassen, O. A., Asherson, P., Banaschewski, T., Bargallo, N., Baumeister, S., Baur-Streubel, R., ... Dima, D. (2022). Cortical thickness across the lifespan: Data from 17,075 healthy individuals aged 3–90 years. *Human Brain Mapping*, *43*(1), 431–451. https://doi.org/https://doi.org/10.1002/hbm.25364

Frid-Adar, M., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). Synthetic data augmentation using gan for improved liver lesion classification. *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, 289–293.

Friston, K. J., Williams, S., Howard, R., Frackowiak, R. S., & Turner, R. (1996). Movement-related effects in fmri time-series. *Magnetic resonance in medicine*, *35*(3), 346–355.

Gazula, H., Tregidgo, H. F. J., Billot, B., Balbastre, Y., William-Ramirez, J., Herisse, R., Casamitjana, A., Melief, E. J., Latimer, C. S., Kilgore, M. D., Montine, M., Robinson, E. D., Blackburn, E., Marshall, M. S., Connors, T. R., Oakley, D. H., Frosch, M. P., Leemput, K. V., Dalca, A. V., ... Iglesias, J. E. (2023). Machine learning of dissection photographs and surface scanning for quantitative 3d neuropathology. *eLife*, *12*.

Georg, M., Preußer, T., & Hahn, H. K. (2010). Global constructive optimization of vascular systems.

Goetz, C. G. (2007). *Textbook of clinical neurology* (Vol. 355). Elsevier Health Sciences.

Gohorbani, A., Natarajan, V., Coz, D. D., & Liu, Y. (2019). Dermgan: Synthetic generation of clinical skin images with pathology.

Gopinath, K., Greve, D. N., Das, S., Arnold, S., Magdamo, C., & Iglesias, J. E. (2023). Cortical analysis of heterogeneous clinical brain mri scans for large-scale neuroimaging studies. *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, 35–45.

Gopinath, K., Greve, D. N., Magdamo, C., Arnold, S., Das, S., Puonti, O., & Iglesias, J. E. (2024). Recon-all-clinical: Cortical surface reconstruction and analysis of heterogeneous clinical brain mri. *arXiv preprint arXiv:2409.03889*.

Gordillo, N., Montseny, E., & Sobrevilla, P. (2013). State of the art survey on MRI brain tumor segmentation. *Magnetic resonance imaging*, *31*(8), 1426–1438.

Graff, C. G. (2016). A new, open-source, multi-modality digital breast phantom. *Medical Imaging 2016: Physics of Medical Imaging*, *9783*, 72–81.

Graves, M. J., & Mitchell, D. G. (2013). Body mri artifacts in clinical practice: A physicist's and radiologist's perspective. *Journal of Magnetic Resonance Imaging*, *38*(2), 269–287.

Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., & Haase, A. (2002). Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, *47*(6), 1202–1210.

Hamarneh, G., & Jassi, P. (2010). Vascusynth: Simulating vascular trees for generating volumetric image data with ground-truth segmentation and tree analysis. *Computerized medical imaging and graphics*, *34*(8), 605–616.

Hewitt, C., Baltruvsaitis, T., Wood, E., Petikam, L., Florentin, L., & Velasquez, H. C. (2023). Procedural humans for computer vision. *arXiv preprint arXiv:2301.01161*.

Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in neural information processing systems*, *33*, 6840–6851.

Hoffmann, M., Abaci Turk, E., Gagoski, B., Morgan, L., Wighton, P., Tisdall, M. D., Reuter, M., Adalsteinsson, E., Grant, P. E., Wald, L. L., & van der Kouwe, A. J. (2021). Rapid head-pose detection for automated slice prescription of fetal-brain MRI. *International Journal of Imaging Systems and Technology*, *31*(3), 1136–1154.

Hoffmann, M., Billot, B., Greve, D. N., Iglesias, J. E., Fischl, B., & Dalca, A. V. (2022). Synthmorph: Learning contrast-invariant registration without acquired images. *IEEE Transactions on Medical Imaging*, *41*(3), 543–558.

Hoffmann, M., Billot, B., Iglesias, J. E., Fischl, B., & Dalca, A. V. (2021). Learning mri contrast-agnostic registration. *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 899–903.

Hoffmann, M., Hoopes, A., Fischl, B., & Dalca, A. V. (2023). Anatomy-specific acquisition-agnostic affine registration learned from fictitious images. *Medical Imaging 2023: Image Processing*, *12464*, 1246402.

Hoffmann, M., Hoopes, A., Greve, D. N., Fischl, B., & Dalca, A. V. (2024). Anatomy-aware and acquisition-agnostic joint registration with SynthMorph. *Imaging Neuroscience*.

Hoopes, A., Hoffmann, M., Greve, D. N., Fischl, B. R., Guttag, J. V., & Dalca, A. V. (2022). Learning the effect of registration hyperparameters with hypermorph. *The journal of machine learning for biomedical imaging*, *1*.

Hoopes, A., Iglesias, J. E., Fischl, B., Greve, D., & Dalca, A. V. (2022). TopoFit: Rapid reconstruction of topologically-correct cortical surfaces. *Proceedings of machine learning research*, *172*, 508.

Hoopes, A., Mora, J. S., Dalca, A. V., Fischl, B., & Hoffmann, M. (2022). SynthStrip: Skull-stripping for any brain image. *NeuroImage*, *260*, 119474.

Iglesias, J. E., Billot, B., Balbastre, Y., Magdamo, C., Arnold, S. E., Das, S., Edlow, B. L., Alexander, D. C., Golland, P., & Fischl, B. (2023). SynthSR: A public AI tool to turn heterogeneous clinical brain scans into high-resolution T1-weighted images for 3D morphometry. *Science advances*, *9*(5), eadd3607.

Iglesias, J. E. (2023). A ready-to-use machine learning tool for symmetric multi-modality registration of brain MRI. *Scientific Reports*, *13*(1), 6657.

Iglesias, J. E., Billot, B., Balbastre, Y., Tabari, A., Conklin, J., Alexander, D. C., Golland, P., Edlow, B. L., & Fischl, B. R. (2020). Joint super-resolution and synthesis of 1 mm isotropic mp-rage volumes from clinical mri exams with scans of different orientation, resolution and contrast. *Neuroimage*, *237*.

Iglesias, J. E., Konukoglu, E., Zikic, D., Glocker, B., Van Leemput, K., & Fischl, B. (2013). Is synthesizing MRI contrast useful for inter-modality analysis? *MICCAI*, 631–38.

Iglesias, J. E., Liu, C.-Y., Thompson, P. M., & Tu, Z. (2011). Robust brain extraction across datasets and comparison with publicly available methods. *IEEE transactions on medical imaging*, *30*(9), 1617–1634.

Iglesias, J. E., Schleicher, R., Laguna, S., Billot, B., Schaefer, P., McKaig, B., Goldstein, J. N., Sheth, K. N., Rosen, M. S., & Kimberly, W. T. (2022). Quantitative brain morphometry of portable low-field-strength MRI using super-resolution machine learning. *Radiology*, *306*(3), e220522.

Jack Jr., C. R., Bernstein, M. A., Fox, N. C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P. J., L. Whitwell, J., Ward, C., Dale, A. M., Felmlee, J. P., Gunter, J. L., Hill, D. L., Killiany, R., Schuff, N., Fox-Bosetti, S., Lin, C., Studholme, C., ... Weiner, M. W. (2008). The alzheimer's disease neuroimaging initiative (adni): Mri methods. *Journal of Magnetic Resonance Imaging*, *27*(4), 685–691. https://doi.org/https://doi.org/10.1002/jmri.21049

Jacobs, L., Mandija, S., Liu, H., van den Berg, C. A., Sbrizzi, A., & Maspero, M. (2023). Generalizable synthetic mri with physics-informed convolutional networks. *Medical Physics*.

Jezzard, P., & Balaban, R. S. (1995). Correction for geometric distortion in echo planar images from b0 field variations. *Magnetic resonance in medicine*, *34*(1), 65–73.

Jiang, J., Hu, Y.-C., Tyagi, N., Zhang, P., Rimner, A., Mageras, G. S., Deasy, J. O., & Veeraraghavan, H. (2018). Tumor-aware, adversarial domain adaptation from ct to mri for lung cancer segmentation. *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part II 11*, 777–785.

Jiang, Y., Chen, H., Loew, M., & Ko, H. (2020). Covid-19 ct image synthesis with a conditional generative adversarial network. *IEEE Journal of Biomedical and Health Informatics*, *25*(2), 441–452.

Jog, A., Hoopes, A., Greve, D. N., Van Leemput, K., & Fischl, B. (2019). PSACNN: Pulse Sequence Adaptive Fast Whole Brain Segmentation. *NeuroImage*, *199*, 553–569.

Johnson, K. A., Fox, N. C., Sperling, R. A., & Klunk, W. E. (2012). Brain imaging in alzheimer disease. *Cold Spring Harbor perspectives in medicine*, *2*(4), a006213.

Jones, D. K., Jones, D. K., Alexander, D. C., Alexander, D. C., Bowtell, R., Cercignani, M., Dell'Acqua, F., McHugh, D. J., McHugh, D. J., Miller, K. L., Palombo, M., Parker, G. J. M., Parker, G. J. M., Rudrapatna, S. U., & Tax, C. M. W. (2018). Microstructural imaging of the human brain with a

'super-scanner': 10 key advantages of ultra-strong gradients for diffusion mri. *NeuroImage*, *182*, 8–38.

Kainz, W., Neufeld, E., Bolch, W. E., Graff, C. G., Kim, C. H., Kuster, N., Lloyd, B. A., Morrison, T. M., Segars, P., Yeom, Y. S., Zankl, M., Xu, X. G., & Tsui, B. M. W. (2019). Advances in computational human phantoms and their applications in biomedical engineering—a topical review. *IEEE Transactions on Radiation and Plasma Medical Sciences*, *3*, 1–23.

Kamnitsas, K., Ledig, C., Newcombe, V. F., Simpson, J. P., Kane, A. D., Menon, D. K., Rueckert, D., & Glocker, B. (2017). Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Medical image analysis*, *36*, 61–78.

Kar, A., Prakash, A., Liu, M.-Y., Cameracci, E., Yuan, J., Rusiniak, M., Acuna, D., Torralba, A., & Fidler, S. (2019). Meta-sim: Learning to generate synthetic datasets. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4551–4560.

Karaçali, B., & Davatzikos, C. (2006). Simulation of Tissue Atrophy Using a Topology Preserving Transformation Model. *IEEE TRANSACTIONS ON MEDICAL IMAGING*, *25*(5).

Karch, R., Neumann, F., Neumann, M., & Schreiner, W. (1999). A three-dimensional model for arterial tree representation, generated by constrained constructive optimization. *Computers in biology and medicine*, *29*(1), 19–38.

Karch, R., Neumann, F., Neumann, M., & Schreiner, W. (2000). Staged growth of optimized arterial model trees. *Annals of biomedical engineering*, *28*, 495–511.

Kazerouni, A., Aghdam, E. K., Heidari, M., Azad, R., Fayyaz, M., Hacihaliloglu, I., & Merhof, D. (2023). Diffusion models in medical imaging: A comprehensive survey. *Medical Image Analysis*, 102846.

Kelley, W., Ngo, N., Dalca, A. V., Fischl, B., Zöllei, L., & Hoffmann, M. (2024). Boosting Skull-Stripping Performance for Pediatric Brain Images. *2024 IEEE 21st International Symposium on Biomedical Imaging (ISBI)*.

Khan, A., Rauf, Z., Sohail, A., Khan, A. R., Asif, H., Asif, A., & Farooq, U. (2023). A survey of the vision transformers and their cnn-transformer based variants. *Artificial Intelligence Review*, *56*(Suppl 3), 2917–2970.

Khanal, B., Ayache, N., & Pennec, X. (2017). Simulating Longitudinal Brain MRIs with Known Volume Changes and Realistic Variations in Image Intensity. *Frontiers in Neuroscience*, *11*, 132.

Kim, S., & Alexander, D. C. (2021). Agcn: Adversarial graph convolutional network for 3d point cloud segmentation. *British Machine Vision Conference*.

Kim, S., Alexander, D. C., Eldaly, A. K., Figini, M., & Tregidgo, H. F. (n.d.). A 3d conditional diffusion model for image quality transfer-an application to low-field mri. *Deep Generative Models for Health Workshop NeurIPS 2023*.

Kingma, D. P., & Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.

Kleesiek, J., Urban, G., Hubert, A., Schwarz, D., Maier-Hein, K., Bendszus, M., & Biller, A. (2016). Deep mri brain extraction: A 3d convolutional neural network for skull stripping. *NeuroImage*, *129*, 460–469.

Kossen, T., Hirzel, M. A., Madai, V. I., Boenisch, F., Hennemuth, A., Hildebrand, K., Pokutta, S., Sharma, K., Hilbert, A., Sobesky, J., Galinovic, I., Khalil, A. A., Fiebach, J. B., & Frey, D. (2022). Toward sharing brain images: Differentially private tof-mra images with segmentation labels using generative adversarial networks. *Frontiers in Artificial Intelligence*, *5*.

Krebs, J., Delingette, H., Mailhé, B., Ayache, N., & Mansi, T. (2019). Learning a probabilistic model for diffeomorphic registration. *IEEE TMI*, *38*(9), 2165–176.

Krebs, J., Mansi, T., Delingette, H., Zhang, L., Ghesu, F. C., Miao, S., Maier, A. K., Ayache, N., Liao, R., & Kamen, A. (2017). Robust non-rigid registration through agent-based action learning. *International Conference on Medical Image Computing and Computer-Assisted Intervention*.

Kundu, A., Li, Y., & Rehg, J. M. (2018). 3d-rcnn: Instance-level 3d object reconstruction via render-and-compare. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3559–3568.

Larson, K. E., & Oguz, I. (2022). Synthetic atrophy for longitudinal cortical surface analyses. *Frontiers in neuroimaging*, *1*, 861687.

LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, *86*(11), 2278–2324.

Li, C.-L., Sohn, K., Yoon, J., & Pfister, T. (2021). Cutpaste: Self-supervised learning for anomaly detection and localization. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9659–9669.

Li, J., Chen, H., Li, Y., Peng, Y., Sun, J., & Pan, P. (2022). Cross-modality synthesis aiding lung tumor segmentation on multi-modal mri images. *Biomedical Signal Processing and Control*, *76*, 103655.

Li, X., Samei, E., Delong, D., Jones, R., Gaca, A., Hollingsworth, C., Maxfield, C., Carrico, C., & Frush, D. (2009). Three-dimensional simulation of lung nodules for paediatric multidetector array ct. *The British journal of radiology*, *82*(977), 401–411.

Liew, S.-L., Anglin, J., Banks, N. W., Sondag, M., Ito, K. L., Kim, H., Chan, J., Ito, J., Jung, C., Khoshab, N., Lefebvre, S., Nakamura, W., Saldana, D., Schmiesing, A., Tran, C., Vo, D., Ard, T., Heydari, P., Kim, B., ... Stroud, A. (2017). A large, open source dataset of stroke anatomical brain images and manual lesion segmentations. *Scientific Data*, *5*.

Lin, C.-H., Kong, C., & Lucey, S. (2018). Learning efficient point cloud generation for dense 3d object reconstruction. *AAAI Conference on Artificial Intelligence (AAAI)*.

Lin, H., Figini, M., D'Arco, F., Ogbole, G., Tanno, R., Blumberg, S. B., Ronan, L., Brown, B. J., Carmichael, D. W., Lagunju, I., Cross, J. H., Fernandez-Reyes, D., & Alexander, D. C. (2023).

Low-field magnetic resonance image enhancement via stochastic image quality transfer. *Medical Image Analysis*, *87*, 102807.

Liu, P., Puonti, O., Hu, X., Alexander, D. C., & Iglesias, J. E. (2023). Brain-id: Learning robust feature representations for brain imaging. *arXiv preprint arXiv:2311.16914*.

Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.

Lorenz, C., Carlsen, I.-C., Buzug, T. M., Fassnacht, C., & Weese, J. (1997). Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2d and 3d medical images. *International Conference on Computer Vision, Virtual Reality, and Robotics in Medicine*, 233–242.

Lustig, M., Donoho, D. L., Santos, J. M., & Pauly, J. M. (2008). Compressed sensing mri. *IEEE signal processing magazine*, *25*(2), 72–82.

Ma, C., Ji, Z., & Gao, M. (2019). Neural style transfer improves 3d cardiovascular mr image segmentation on inconsistent data. *International Conference on Medical Image Computing and Computer-Assisted Intervention*.

Ma, D., Gulani, V., Seiberlich, N., Liu, K., Sunshine, J. L., Duerk, J. L., & Griswold, M. A. (2013). Magnetic resonance fingerprinting. *Nature*, *495*(7440), 187–192.

Ma, Q., Li, L., Robinson, E. C., Kainz, B., Rueckert, D., & Alansary, A. (2022). CortexODE: Learning cortical surface reconstruction by neural ODEs. *IEEE Transactions on Medical Imaging*, *42*(2), 430–443.

Ma, Q., Robinson, E. C., Kainz, B., Rueckert, D., & Alansary, A. (2021). PialNN: A fast deep learning framework for cortical pial surface reconstruction. *Machine Learning in Clinical Neuroimaging: 4th International Workshop, MLCN 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 4*, 73–81.

Malik, S. J., Kenny, G. D., & Hajnal, J. V. (2011). Slice profile correction for transmit sensitivity mapping using actual flip angle imaging. *Magnetic resonance in medicine*, *65*(5), 1393–1399.

Marques, J. P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P.-F., & Gruetter, R. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *Neuroimage*, *49*(2), 1271–1281.

Matthews, P. M., Honey, G. D., & Bullmore, E. T. (2006). Applications of fMRI in translational medicine and clinical practice. *Nature Reviews Neuroscience*, *7*(9), 732–744.

Menten, M. J., Paetzold, J. C., Dima, A., Menze, B. H., Knier, B., & Rueckert, D. (2022). Physiology-based simulation of the retinal vasculature enables annotation-free segmentation of oct angiographs. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 330–340.

Menze, B. H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.-A., Arbel, T., Avants, B. B., Ayache, N., Buendia, P., Collins, D. L., Cordier, N., ... Van Leemput, K. (2015). The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging*, *34*(10), 1993–2024. https://doi.org/10.1109/TMI.2014.2377694

Milletari, F., Navab, N., & Ahmadi, S.-A. (2016). V-net: Fully convolutional neural networks for volumetric medical image segmentation. *2016 fourth international conference on 3D vision (3DV)*, 565–571.

Modat, M., et al. (2010). Fast free-form deformation using graphics processing units. *Comput Meth Prog Bio*, *98*(3), 278–84.

Monje, M. (2018). Myelin plasticity and nervous system function. *Annual review of neuroscience*, *41*, 61–76.

Montoya, A., Price, B. H., Menear, M., & Lepage, M. (2006). Brain imaging and cognitive dysfunctions in huntington's disease. *Journal of Psychiatry and Neuroscience*, *31*(1), 21–29.

Neumann, F., Schreiner, W., & Neumann, M. (1995). Constrained constructive optimization of binary branching arterial tree models. *WIT Transactions on The Built Environment*, *14*.

Nieuwenhuys, R., & Broere, C. A. (2017). A map of the human neocortex showing the estimated overall myelin content of the individual architectonic areas based on the studies of adolf hopf. *Brain Structure and Function*, *222*, 465–480.

Nikolenko, S. I. (2021). *Synthetic data for deep learning* (Vol. 174). Springer.

Nolan, A. L., Petersen, C., Iacono, D., Mac Donald, C. L., Mukherjee, P., Van Der Kouwe, A., Jain, S., Stevens, A., Diamond, B. R., Wang, R., Markowitz, A. J., Fischl, B., Perl, D. P., Manley, G. T., Keene, C. D., Diaz-Arrastia, R., Edlow, B. L., the TRACK-TBI Investigators, Adeoye, O., ... Zafonte, R. (2021). Tractography-Pathology Correlations in Traumatic Brain Injury: A TRACK-TBI Study. *Journal of Neurotrauma*, *38*(12), 1620–1631. https://doi.org/10.1089/neu.2020.7373

Nuesch, E., Trelle, S., Reichenbach, S., Rutjes, A. W., Bürgi, E., Scherer, M., Altman, D. G., & Jüni, P. (2009). The effects of excluding patients from the analysis in randomised controlled trials: Meta-epidemiological study. *Bmj*, *339*.

Ouyang, C., Chen, C., Li, S., Li, Z., Qin, C., Bai, W., & Rueckert, D. (2022). Causality-inspired single-source domain generalization for medical image segmentation. *IEEE Transactions on Medical Imaging*, *42*(4), 1095–1106.

Patenaude, B., Smith, S. M., Kennedy, D. N., & Jenkinson, M. (2011). A bayesian model of shape and appearance for subcortical brain segmentation. *Neuroimage*, *56*(3), 907–922.

Pereira, S., Pinto, A., Alves, V., & Silva, C. A. (2016). Brain tumor segmentation using convolutional neural networks in MRI images. *IEEE transactions on medical imaging*, *35*(5), 1240–1251.

Pezeshk, A., Sahiner, B., Zeng, R., Wunderlich, A., Chen, W., & Petrick, N. (2015). Seamless insertion of pulmonary nodules in chest ct images. *IEEE Transactions on Biomedical Engineering*, *62*(12), 2812–2827.

Pinaya, W. H., Tudosiu, P.-D., Dafflon, J., Da Costa, P. F., Fernandez, V., Nachev, P., Ourselin, S., & Cardoso, M. J. (2022). Brain imaging generation with latent diffusion models. *MICCAI Workshop on Deep Generative Models*, 117–126.

Politis, M. (2014). Neuroimaging in parkinson disease: From research setting to clinical practice. *Nature Reviews Neurology*, *10*(12), 708–722.

Prados, F., Cardoso, M. J., Cawley, N., Kanber, B., Ciccarelli, O., Wheeler-Kingshott, C. A. G., & Ourselin, S. (2016). Fully automated patch-based image restoration: Application to pathology inpainting. *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: Second International Workshop, BrainLes 2016, with the Challenges on BRATS, ISLES and mTOP 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 17, 2016, Revised Selected Papers 2*, 3–15.

Prastawa, M., Bullitt, E., & Gerig, G. (2009). Simulation of brain tumors in MR images for evaluation of segmentation efficacy. *Medical Image Analysis*, *13*(2), 297–311.

Pruessmann, K. P., Weiger, M., Scheidegger, M. B., & Boesiger, P. (1999). Sense: Sensitivity encoding for fast mri. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, *42*(5), 952–962.

Rinck, P. A. (2019). *Magnetic resonance in medicine: A critical introduction*. BoD–Books on Demand.

Rohé, M., Datar, M., Heimann, T., Sermesant, M., & Pennec, X. (2017). SVF-Net: Learning deformable image registration using shape matching. *MICCAI*, 266–74.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Ros, G., Sellart, L., Materzynska, J., Vazquez, D., & Lopez, A. M. (2016). The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3234–3243.

Roy, S., Butman, J. A., & Pham, D. L. (2017). Robust skull stripping using multiple mr image contrasts insensitive to pathology. *NeuroImage*, *146*, 132–147.

Rueckert, D., Sonoda, L. I., Hayes, C., Hill, D. L., Leach, M. O., & Hawkes, D. J. (1999). Nonrigid registration using free-form deformations: Application to breast mr images. *IEEE TMI*, *18*(8), 712–21.

Rusak, F., Cruz, R. S., Lebrat, L., Hlinka, O., Fripp, J., Smith, E., Fookes, C., Bradley, A. P., & Bourgeat, P. T. (2022). Quantifiable brain atrophy synthesis for benchmarking of cortical thickness estimation methods. *Medical image analysis*, *82*, 102576.

Sadeghi, F., & Levine, S. (2016). Cad2rl: Real single-image flight without a single real image. *ArXiv*, *abs/1611.04201*.

Salehi, S. S. M., Erdogmus, D., & Gholipour, A. (2017). Auto-context convolutional neural network (auto-net) for brain extraction in magnetic resonance imaging. *IEEE transactions on medical imaging*, *36*(11), 2319–2330.

Salem, M., Valverde, S., Cabezas, M., Pareto, D., Oliver, A., Salvi, J., Rovira, À., & Lladó, X. (2019). Multiple sclerosis lesion synthesis in mri using an encoder-decoder u-net. *IEEE Access*, *7*, 25171–25184.

Sarno, A., Mettivier, G., di Franco, F., Varallo, A., Bliznakova, K., Hernandez, A. M., Boone, J. M., & Russo, P. (2021). Dataset of patient-derived digital breast phantoms for in silico studies in breast computed tomography, digital breast tomosynthesis, and digital mammography. *Medical Physics*, *48*(5), 2682–2693.

Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., & Kikinis, R. (1998). Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical image analysis*, *2*(2), 143–168.

Sauer, T. J., Bejan, A., Segars, P., & Samei, E. (2023). Development and ct image-domain validation of a computational lung lesion model for use in virtual imaging trials. *Medical physics*, *50*(7), 4366–4378.

Schneider, M., Reichold, J., Weber, B., Székely, G., & Hirsch, S. (2012). Tissue metabolism driven arterial tree generation. *Medical image analysis*, *16*(7), 1397–1414.

Ségonne, F., Dale, A. M., Busa, E., Glessner, M., Salat, D., Hahn, H. K., & Fischl, B. (2004). A hybrid approach to the skull stripping problem in mri. *Neuroimage*, *22*(3), 1060–1075.

Sengupta, A., Sharma, D., & Badano, A. (2021). Computational model of tumor growth for in silico trials. *Medical Imaging 2021: Physics of Medical Imaging*, *11595*, 1262–1270.

Shah, S., Dey, D., Lovett, C., & Kapoor, A. (2018). Airsim: High-fidelity visual and physical simulation for autonomous vehicles. *Field and Service Robotics: Results of the 11th International Conference*, 621–635.

Shang, Z., Turja, M. A., Feczko, E., Houghton, A., Rueter, A. R., Moore, L. A., Snider, K., Hendrickson, T., Reiners, P., Stoyell, S. M., Kardan, O., Elison, J. T., & Styner, M. (2022). Learning strategies for contrast-agnostic segmentation via synthseg for infant mri data. *Proceedings of machine learning research*, *172*, 1075–1084.

Shin, H.-C., Tenenholtz, N. A., Rogers, J. K., Schwarz, C. G., Senjem, M. L., Gunter, J. L., Andriole, K. P., & Michalski, M. (2018). Medical image synthesis for data augmentation and anonymization using generative adversarial networks. *Simulation and Synthesis in Medical Imaging: Third International Workshop, SASHIMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 3*, 1–11.

Shin, S., Yun, S. D., & Shah, N. J. (2023). T2* quantification using multi-echo gradient echo sequences: A comparative study of different readout gradients. *Scientific Reports*, *13*(1), 1138.

Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of big data*, *6*(1), 1–48.

Sijbers, J., den Dekker, A. J., Raman, E., & Van Dyck, D. (1999). Parameter estimation from magnitude mr images. *International Journal of imaging systems and technology*, *10*(2), 109–114.

Simon, R. P., Aminoff, M. J., & Greenberg, D. A. (2009). *Clinical neurology* (Vol. 20). Lange Medical Books/McGraw-Hill.

Sizikova, E., Saharkhiz, N., Sharma, D., Lago, M., Sahiner, B., Delfino, J., & Badano, A. (2024). Knowledge-based in silico models and dataset for the comparative evaluation of mammography ai for a range of breast characteristics, lesion conspicuities and doses. *Advances in Neural Information Processing Systems*, *36*.

Smith, A. D. C., Crum, W. R., Hill, D. L. G., Thacker, N. A., & Bromiley, P. A. (2003). Biomechanical simulation of atrophy in MR images. *Medical Imaging 2003: Image Processing*, *5032*, 481–490.

Smith, S. M. (2002). Fast robust automated brain extraction. *Human brain mapping*, *17*(3), 143–155.

Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E., Johansen-Berg, H., Bannister, P. R., Luca, M. D., Drobnjak, I., Flitney, D. E., Niazy, R. K., Saunders, J., Vickers, J., Zhang, Y., Stefano, N. D., Brady, J. M., & Matthews, P. M. (2004). Advances in functional and structural mr image analysis and implementation as fsl [Mathematics in Brain Imaging]. *NeuroImage*, *23*, S208–S219. https://doi.org/https://doi.org/10.1016/j.neuroimage.2004.07.051

Sun, H., Plawinski, J., Subramaniam, S., Jamaludin, A., Kadir, T., Readie, A., Ligozio, G., Ohlssen, D., Baillie, M., & Coroller, T. (2023). A deep learning approach to private data sharing of medical images using conditional generative adversarial networks (gans). *Plos one*, *18*(7), e0280316.

Szczerba, D., & Székely, G. (2005). Simulating vascular systems in arbitrary anatomies. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005: 8th International Conference, Palm Springs, CA, USA, October 26-29, 2005, Proceedings, Part II 8*, 641–648.

Tanno, R., Worrall, D. E., Kaden, E., & Alexander, D. C. (2020). Uncertainty modelling in deep learning for safer neuroimage enhancement: Demonstration in diffusion mri. *NeuroImage*, *225*.

Taye, M. M. (2023). Understanding of machine learning with deep learning: Architectures, workflow, applications and future directions. *Computers*, *12*(5), 91.

Teixeira, B., Singh, V., Chen, T., Ma, K., Tamersoy, B., Wu, Y., Balashova, E., & Comaniciu, D. (2018). Generating synthetic x-ray images of a person from the surface geometry. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 9059–9067.

Thambawita, V., Salehi, P., Sheshkal, S. A., Hicks, S. A., Hammer, H. L., Parasa, S., Lange, T. d., Halvorsen, P., & Riegler, M. A. (2022). Singan-seg: Synthetic training data generation for medical image segmentation. *PloS one*, *17*(5), e0267976.

The Alzheimer's Disease Neuroimaging Initiative, The CHARGE Consortium, EPIGEN, IMAGEN, SYS, Hibar, D. P., Stein, J. L., Renteria, M. E., Arias-Vasquez, A., Desrivières, S., Jahanshad, N., Toro, R., Wittfeld, K., Abramovic, L., Andersson, M., Aribisala, B. S., Armstrong, N. J., Bernard, M., Bohlken, M. M., ... Medland, S. E. (2015). Common genetic variants influence human subcortical brain structures. *Nature*, *520*(7546), 224–229. https://doi.org/10.1038/nature14101

Thompson, P. M., Jahanshad, N., Ching, C. R. K., Salminen, L. E., Thomopoulos, S. I., Bright, J., Baune, B. T., Bertolín, S., Bralten, J., Bruin, W. B., Bülow, R., Chen, J., Chye, Y., Dannlowski, U., De Kovel, C. G. F., Donohoe, G., Eyler, L. T., Faraone, S. V., Favre, P., ... for the ENIGMA Consortium. (2020). ENIGMA and global neuroscience: A decade of large-scale studies of the brain in health and disease across more than 40 countries. *Translational Psychiatry*, *10*(1), 100. https://doi.org/10.1038/s41398-020-0705-1

Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., & Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 23–30.

Todorov, E., Erez, T., & Tassa, Y. (2012). Mujoco: A physics engine for model-based control. *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033.

Tregidgo, H. F. J., Casamitjana, A., Latimer, C. S., Kilgore, M. D., Robinson, E., Blackburn, E., Van Leemput, K., Fischl, B., Dalca, A. V., Donald, C. L. M., Keene, C. D., & Iglesias, J. E. (2020). 3d reconstruction and segmentation of dissection photographs for mri-free neuropathology. In A. L. Martel, P. Abolmaesumi, D. Stoyanov, D. Mateus, M. A. Zuluaga, S. K. Zhou, D. Racoceanu, & L. Joskowicz (Eds.), *Medical image computing and computer assisted intervention – miccai 2020* (pp. 204–214). Springer International Publishing.

Tremblay, J., Prakash, A., Acuna, D., Brophy, M., Jampani, V., Anil, C., To, T., Cameracci, E., Boochoon, S., & Birchfield, S. (2018). Training deep networks with synthetic data: Bridging the reality gap by domain randomization. *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 969–977.

Van Essen, D. C., & Glasser, M. F. (2014). In vivo architectonics: A cortico-centric perspective. *Neuroimage*, *93*, 157–164.

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., & Ugurbil, K. (2013). The wu-minn human connectome project: An overview [Mapping the Connectome]. *NeuroImage*, *80*, 62–79. https://doi.org/https://doi.org/10.1016/j.neuroimage.2013.05.041

Van Leemput, K., Maes, F., Vandermeulen, D., Colchester, A., & Suetens, P. (2001). Automated segmentation of multiple sclerosis lesions by model outlier detection. *IEEE transactions on medical imaging*, *20*(8), 677–688.

Varadarajan, D., Bouman, K. L., van der Kouwe, A., Fischl, B., & Dalca, A. V. (2021). Unsupervised learning of mri tissue properties using mri physics models. *arXiv preprint arXiv:2107.02704*.

Varadarajan, D., & Haldar, J. P. (2015). A majorize-minimize framework for rician and non-central chi mr images. *IEEE transactions on medical imaging*, *34*(10), 2191–2202.

Varol, G., Romero, J., Martin, X., Mahmood, N., Black, M. J., Laptev, I., & Schmid, C. (2017). Learning from synthetic humans. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 109–117.

Waheed, A., Goyal, M., Gupta, D., Khanna, A., Al-turjman, F. M., & Pinheiro, P. R. (2020). Covidgan: Data augmentation using auxiliary classifier gan for improved covid-19 detection. *IEEE Access*, *8*, 91916–91923.

Wang, M., & Deng, W. (2018). Deep visual domain adaptation: A survey. *Neurocomputing*, *312*, 135–153.

Webster, J. M., Grabowski, T. J., Madhyastha, T. M., Gibbons, L. E., Keene, C. D., & Latimer, C. S. (2021). Leveraging neuroimaging tools to assess precision and accuracy in an alzheimer's disease neuropathologic sampling protocol. *Frontiers in Neuroscience*, *15*.

Wilms, M., Bannister, J. J., Mouches, P., MacDonald, M. E., Rajashekar, D., Langner, S., & Forkert, N. D. (2022). Invertible modeling of bidirectional relationships in neuroimaging with normalizing flows: Application to brain aging. *IEEE Transactions on Medical Imaging*, *41*(9), 2331–2347.

Wong, H. E., Rakic, M., Guttag, J., & Dalca, A. V. (2023). Scribbleprompt: Fast and flexible interactive segmentation for any medical image. *arXiv:2312.07381*.

Wood, E., Baltruvsaitis, T., Hewitt, C., Dziadzio, S., Cashman, T. J., & Shotton, J. (2021). Fake it till you make it: Face analysis in the wild using synthetic data alone. *Proceedings of the IEEE/CVF international conference on computer vision*, 3681–3691.

Wood, E., Baltruvsaitis, T., Hewitt, C., Johnson, M., Shen, J., Milosavljevic, N., Wilde, D., Garbin, S. J., Raman, C., Sharp, T., Stojiljković, I. N., Cashman, T., & Valentin, J. P. C. (2022). 3d face reconstruction with dense landmarks. *European Conference on Computer Vision*.

Wu, B., Wan, A., Yue, X., & Keutzer, K. (2017). Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 1887–1893.

Xanthis, C. G., Filos, D., Haris, K., & Aletras, A. H. (2021). Simulator-generated training datasets as an alternative to using patient data for machine learning: An example in myocardial segmentation with mri. *Computer Methods and Programs in Biomedicine*, *198*, 105817.

Xu, Z., Liu, D., Yang, J., Raffel, C., & Niethammer, M. (2020). Robust and generalizable visual representation learning via random convolutions. *arXiv preprint arXiv:2007.13003*.

Yamashita, R., Long, J., Banda, S., Shen, J., & Rubin, D. (2021). Learning domain-agnostic visual representation for computational pathology using medically-irrelevant style transfer augmentation. *IEEE Transactions on Medical Imaging*, *40*, 3945–3954.

Yang, X., Kwitt, R., Styner, M., & Niethammer, M. (2017). Quicksilver: Fast predictive image registration–a deep learning approach. *NeuroImage*, *158*, 378–96.

Young, S. I., Balbastre, Y., Dalca, A. V., Wells, W. M., Iglesias, J. E., & Fischl, B. (2022). Superwarp: Supervised learning and warping on u-net for invariant subvoxel-precise registration. *International Workshop on Biomedical Image Registration*, 103–115.

Young, S. I., Dalca, A. V., Ferrante, E., Golland, P., Metzler, C. A., Fischl, B., & Iglesias, J. E. (2023). Supervision by denoising. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Zaitsev, M., Maclaren, J., & Herbst, M. (2015). Motion artifacts in mri: A complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging*, *42*(4), 887–901.

Zhang, G., Cui, K., Hung, T.-Y., & Lu, S. (2021). Defect-gan: High-fidelity defect synthesis for automated defect inspection. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2523–2533.

Zhang, H., Bakshi, R., Bagnato, F., & Oguz, I. (2020). Robust multiple sclerosis lesion inpainting with edge prior. *Machine Learning in Medical Imaging*, *12436*, 120–129.

Zhao, A., Balakrishnan, G., Durand, F., Guttag, J. V., & Dalca, A. V. (2019). Data augmentation using learned transformations for one-shot medical image segmentation. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8543–8553.

Zhao, C., Dewey, B. E., Pham, D. L., Calabresi, P. A., Reich, D. S., & Prince, J. L. (2020). Smore: A self-supervised anti-aliasing and super-resolution algorithm for mri using deep learning. *IEEE transactions on medical imaging*, *40*(3), 805–817.

Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R., & Rosen, M. S. (2018). Image reconstruction by domain-transform manifold learning. *Nature*, *555*(7697), 487–492.

Zunair, H., & Hamza, A. B. (2021). Synthesis of covid-19 chest x-rays using unpaired image-to-image translation. *Social network analysis and mining*, *11*, 1–12.