


# Unified 3D MRI Representations via Sequence-Invariant Contrastive Learning

Liam Chalcraft<sup>1</sup> 

L.CHALCROFT@CS.UCL.AC.UK

<sup>1</sup> Department of Imaging Neuroscience, University College London, UK

Jenny Crinion<sup>2</sup> 

<sup>2</sup> Institute of Cognitive Neuroscience, University College London, UK

Cathy J. Price<sup>1</sup>

John Ashburner<sup>1</sup> 

**Editors:** Under Review for MIDL 2025

## Abstract

Self-supervised deep learning has accelerated 2D natural image analysis but remains difficult to translate into 3D MRI, where data are scarce and pre-trained 2D backbones cannot capture volumetric context. We present a *sequence-invariant* self-supervised framework leveraging quantitative MRI (qMRI). By simulating multiple MRI contrasts from a single 3D qMRI scan and enforcing consistent representations across these contrasts, we learn anatomy-centric rather than sequence-specific features. This yields a robust 3D encoder that performs strongly across varied tasks and protocols. Experiments on healthy brain segmentation (IXI), stroke lesion segmentation (ARC), and MRI denoising show significant gains over baseline SSL approaches, especially in low-data settings (up to +8.3% Dice, +4.2 dB PSNR). Our model also generalises effectively to unseen sites, demonstrating potential for more scalable and clinically reliable volumetric analysis. All code and trained models are publicly available.

## 1. Introduction

Deep learning has advanced medical image analysis across a range of tasks, from image registration (Balakrishnan et al., 2019) to semantic segmentation (Chalcraft et al., 2023) and visual question-answering (Narayanan et al., 2024). However, unique challenges arise when working with 3D MRI data, including increased computational demands and the difficulty of applying 2D pre-trained models to volumetric contexts (Ma et al., 2024). Although large-scale 3D datasets (Li et al., 2024) and models (Wu et al., 2024) have recently emerged, fine-tuning them for specific clinical tasks remains non-trivial due to inevitable domain shifts (Yang et al., 2024).

Self-supervised learning (SSL) offers a promising means of learning robust representations without the need for large labelled datasets (Chaitanya et al., 2020; Huang et al., 2023). Yet, existing SSL methods often treat each MRI sequence as a separate domain, neglecting the shared anatomical information across contrast variations. In contrast, we leverage the observation that different MRI sequences, despite their unique contrast properties, encode the same underlying anatomy. By incorporating physics-based knowledge of MRI signal generation into an SSL framework, we can learn representations that are robust to sequence changes while preserving anatomical fidelity.

In this paper, we make the following contributions:

- We propose a sequence-invariant self-supervised learning framework based on physics-driven contrast simulation, enabling robust 3D MRI representations across multiple sequences.
- We show significant performance gains in low-data regimes, with up to 8.3% Dice improvement in segmentation and a 4.2 dB PSNR increase for denoising when using just 1% of the training data.
- We demonstrate enhanced cross-sequence generalisation and site adaptation from a single pre-trained encoder, particularly in out-of-domain scenarios.
- We provide comprehensive evaluations on three diverse tasks - healthy brain segmentation, stroke lesion segmentation, and image denoising - highlighting the clinical utility of our approach.

Our method addresses key problems in medical imaging by enabling robust feature learning across different sites and sequences, even with limited annotated data. This work takes a step towards developing more generalisable and clinically applicable models. We release all code and backbone weights at [github.com/liamchalcroft/contrast-squared](https://github.com/liamchalcroft/contrast-squared).

## 2. Related Work

We briefly review three core areas that underpin this work: contrastive learning, robust representations in 3D medical imaging, and quantitative MRI (qMRI).

### 2.1. Contrastive Learning

Self-supervised learning (SSL) can leverage unlabelled data by creating proxy tasks that encourage useful invariances in learned representations. Techniques include predictive coding (van den Oord et al., 2019; Hénaff et al., 2020), masked image modelling (He et al., 2021; Xie et al., 2022), and contrastive learning (Chen et al., 2020; He et al., 2020).

Recent contrastive methods such as SimCLR (Chen et al., 2020) and MoCo (He et al., 2020) learn representations by aligning features from different augmented views, while BYOL (Grill et al., 2020) and Barlow Twins (Zbontar et al., 2021) reduce the reliance on explicit negative samples or introduce redundancy reduction.

### 2.2. Robust Representations in Medical Imaging

Clinical MRI segmentation tasks face challenges when transferring models to new hospitals or protocols. Public benchmarks often involve a small set of consistent sequences, limiting models to scenarios where training and testing domains match (e.g. T1w-only). Real-world deployment must handle diverse sequences and acquisition conditions.

Existing domain adaptation methods typically require multiple unlabelled images (Wang et al., 2021) or prior knowledge of the target domain (Dorent et al., 2023), which is not always feasible. SynthSeg (Billot et al., 2023) addresses this by randomising tissue contrast with synthetic data, with subsequent work showing the transferrability of the learned representations to new tasks (Liu et al., 2024a,b). The performance of these methods is expected to be constrained by the quality of the synthetic data, which may fail to capture specific

anatomical details desired in downstream tasks. Similarly, Meyer et al. (2021) adjust contrast on specific regions in real images, but this approach is restricted to modest in-domain variations rather than full sequence simulation. Meanwhile, Roschewitz et al. (2024) demonstrate that generating counterfactual views can boost domain robustness for 2D chest X-ray encoders. We extend these insights to 3D MRI for sequence-invariant representations.

### 2.3. Quantitative MRI

Quantitative MRI (qMRI) acquires per-voxel parameter maps (*e.g.*,  $R_1$ ,  $R_2^*$ ) that govern the signal formation in conventional scans (Weiskopf et al., 2021). These maps facilitate the simulation of numerous synthetic MRI sequences from a single qMRI acquisition (Tanenbaum et al., 2017), improving model robustness under domain shift. For example, synthesised multi-contrast data has led to enhanced results in healthy brain parcellation (Borges et al., 2021) and pathology segmentation (Chalcroft et al., 2024). Other methods rely on MR fingerprinting (Ma et al., 2013) to derive similar quantitative maps (Adams et al., 2024), further expanding opportunities for sequence-invariant learning.

## 3. Methods

We propose a sequence-invariant self-supervised learning (SSL) approach for robust 3D MRI representations. Figure 1 gives an overview of the method, which comprises:

1. A backbone encoder trained via contrastive learning on augmented volume pairs;
2. A decoder to reconstruct the original images, promoting semantic consistency for dense tasks;
3. A physics-based generative model to simulate multiple MRI sequences from quantitative MRI (qMRI) parameter maps, ensuring sequence invariance.

### 3.1. Self-Supervised Learning

We adopt SimCLR (Chen et al., 2020) as our core contrastive framework, though other SSL methods could also be used (Zbontar et al., 2021; Bardes et al., 2022; Oquab et al., 2024). Following Tang et al. (2022), we incorporate an additional reconstruction branch. Specifically:

- **Contrastive branch:** We create two augmented 3D views of a single input volume. Each view is passed through a shared encoder, producing latent vectors  $(z_i, z_j)$ . A contrastive loss encourages  $z_i$  and  $z_j$  to be similar while remaining distinct from other samples in the batch. This step induces a rich feature representation that generalises well across domains.
- **Reconstruction branch:** A lightweight decoder reconstructs the original volume from the latent features after removing artificially added artefacts (*e.g.*, noise, dropout). An  $L_1$  loss enforces pixel-level fidelity.

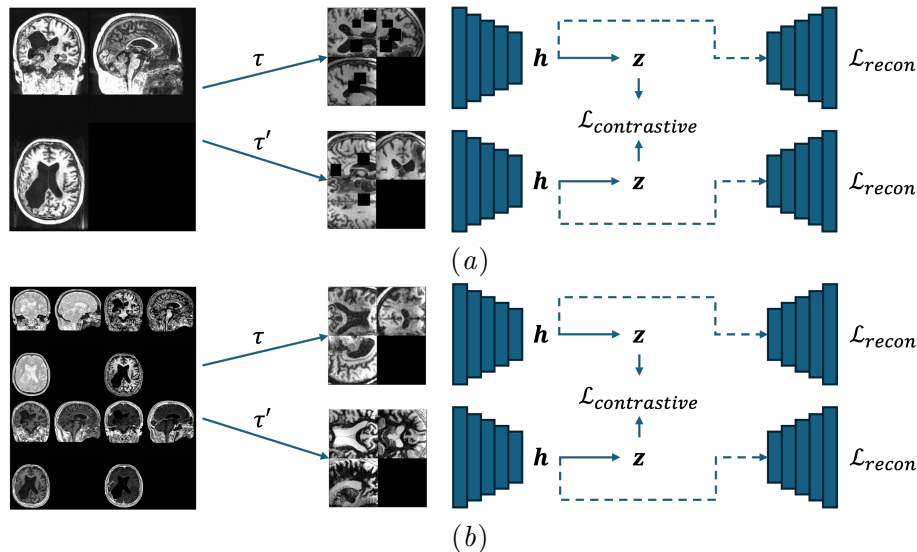


Figure 1: **Overview of the proposed SSL approach.** (1(a)) An MPRAGE volume is augmented into two random views. We extract a feature vector  $h$  via the backbone encoder, project it to  $z$  for a contrastive loss  $\mathcal{L}_{contrastive}$ , and use a decoder to optimise a reconstruction/inpainting loss  $\mathcal{L}_{recon}$ . (1(b)) We generalise this by simulating multiple scanner sequences from qMRI parameter maps, enabling sequence-invariant learning.

For spatial augmentations, each volume is randomly cropped, rotated, sheared, and flipped. We then apply MRI-specific augmentations such as non-uniform intensity fields (Van Leemput et al., 1999), Gibbs artefacts (Morelli et al., 2011), Rician noise (Gudbjartsson and Patz, 1995), and random cuboid dropout (Pathak et al., 2016). In the baseline version, we generate these augmented views from simulated MPRAGE images. In our sequence-invariant framework, we instead use parameter maps to simulate diverse MRI sequences (Sec. 3.2), enabling the encoder to learn anatomy-centric features rather than sequence-specific contrast. We train a model **SeqAug** that generates two views from a single simulated sequence, and a second model **SeqInv** that generates the two views from distinct sequence simulations, formally encouraging invariance to choice of MRI sequence.

### 3.2. Physics-Based Generative Model

We leverage qMRI maps (PD,  $R_1$ ,  $R_2^*$ , MT) to synthesise multiple MRI contrasts from a single subject. Each voxel’s tissue parameters are passed through forward models approximating various standard MRI sequences (FSE, GRE, FLAIR, MPRAGE). Full signal equations are derived from known relaxation properties (Appendix A), and Rician noise is added for realism. By sampling different scanner parameters (*e.g.*, echo time, flip angle), we obtain a range of synthetic images sharing identical anatomical structure but differing in appearance. All simulations use the NITorch library.\*

\*<https://github.com/balbasty/nitorch>

## 4. Experiments and Results

We evaluate our sequence-invariant approach on three downstream tasks: healthy brain segmentation, stroke lesion segmentation, and MRI denoising. Following standard practice, we measure segmentation performance using the Dice Similarity Coefficient (DSC) and 95th percentile Hausdorff Distance (HD95), and denoising performance using Peak Signal-to-Noise Ratio (PSNR).

### 4.1. Implementation Details and Data Setup

**Pretraining.** We train a 3D CNN encoder (architecture detailed in Appendix B) using three different pretraining strategies:

- **Base:** Trained on real MPRAGE scans only, using standard contrastive augmentations
- **SeqAug:** Trained on simulated sequence images, generating two views from the same sequence
- **SeqInv:** Generating two views from distinct sequence simulations, explicitly enforcing sequence invariance

All models use the NT-Xent loss (Chen et al., 2020) with temperature  $\tau = 0.5$  and an optional reconstruction loss with equal weighting. The pretraining dataset consists of 51 qMRI volumes (22 healthy, 29 stroke subjects), with sequence simulation performed using Bloch equations for **SeqAug** and **SeqInv**.

**Downstream Tasks.** Once pretraining is complete, we freeze the encoder and optimise a U-Net decoder for:

- **Healthy Brain Segmentation:** T1w, T2w, and PDw volumes from the IXI dataset (Robinson et al., 2010), segmented into background, grey matter, white matter, and CSF. We train on  $96^3$  patches with affine and intensity augmentations, using a combined Dice + cross-entropy loss. For training, a maximum of 226 subjects are available from the GST site, with 31 reserved for validation and a further 65 for the in-domain test set. For out-of-domain testing, there are 185 and 74 subjects available in the HH and IOP sites respectively.
- **Stroke Lesion Segmentation:** T1w, T2w, and FLAIR from the ARC dataset (Gibson et al., 2024). Lesions are often small, so we employ higher class weighting. We use  $96^3$  patches and the same augmentations, optimising a combined Dice + cross-entropy loss. The T1w, T2w and FLAIR sequences are distributed in respective train/validation/test splits of (142/20/41), (159/22/47) and (59/8/18).
- **MRI Denoising:** We add synthetic noise ( $\sigma = 0.2$ ) to clean IXI scans normalised to a zero mean and unit standard deviation. The network predicts the noise, which is subtracted from the input to produce the denoised image. We evaluate the result via PSNR on the same IXI splits used for healthy segmentation.

All models use  $96^3$  patches with standard augmentations. Training details including optimization strategy, learning rate schedules, and batch sizes are provided in Appendix B.

## 4.2. Evaluation Metrics

**Peak Signal-to-Noise Ratio (PSNR)** Assesses image quality by comparing the maximum possible signal with the noise. For an image with maximum pixel value  $L$ ,  $\text{PSNR} = 20 \cdot \log_{10}\left(\frac{L}{\sqrt{\text{MSE}}}\right)$ , where  $\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$  measures the average error between predicted and ground truth images.

**Dice Similarity Coefficient (DSC)** Measures overlap  $\text{DSC}(Y, \hat{Y}) = \frac{2|Y \cap \hat{Y}|}{|Y| + |\hat{Y}|}$  between a predicted segmentation  $\hat{Y}$  and ground truth  $Y$ . Values range from 0 (no overlap) to 1 (perfect overlap).

**95% Hausdorff Distance (HD95)** Reflects boundary accuracy by measuring the 95th percentile of all directed distances between segmentation boundaries. Smaller values indicate better delineation of anatomical edges.

## 4.3. Quantitative Results

### 4.3.1. HEALTHY BRAIN SEGMENTATION

Table 1 compares DSC scores for T1w, T2w, and PDw images from the IXI dataset, with varying training data proportions (1%, 10%, 100% of the total available data). Our sequence-invariant (**SeqInv**) model consistently outperforms the baseline (**Base**), especially in low-data settings and out-of-domain sites (HH, IOP). Meanwhile, the sequence-augmented (**SeqAug**) model provides moderate gains, particularly on T2w.

Table 1: Healthy brain tissue segmentation performance using Dice Similarity Coefficient (higher is better). Values show mean  $\pm$  std, with **bold** and underlined indicating best and second-best results. GST represents the training domain.

|                      | 1% Training Data                  |                                   |                                   | 10% Training Data                 |                                   |                                   | 100% Training Data                |                                   |                                   |
|----------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
|                      | Base                              | SeqAug                            | SeqInv                            | Base                              | SeqAug                            | SeqInv                            | Base                              | SeqAug                            | SeqInv                            |
| <b>In Domain</b>     |                                   |                                   |                                   |                                   |                                   |                                   |                                   |                                   |                                   |
| GST [T1w]            | <u>55.1 <math>\pm</math> 11.2</u> | 38.5 $\pm$ 12.0                   | <b>56.0 <math>\pm</math> 12.0</b> | <b>69.3 <math>\pm</math> 8.2</b>  | 67.2 $\pm$ 10.2                   | <u>67.9 <math>\pm</math> 9.8</u>  | <b>89.6 <math>\pm</math> 4.2</b>  | 84.1 $\pm$ 8.5                    | <u>85.5 <math>\pm</math> 6.5</u>  |
| GST [T2w]            | <b>65.4 <math>\pm</math> 6.3</b>  | <u>56.9 <math>\pm</math> 7.6</u>  | 47.7 $\pm$ 11.2                   | <b>84.2 <math>\pm</math> 4.7</b>  | <u>79.0 <math>\pm</math> 4.8</u>  | 68.6 $\pm$ 8.6                    | <u>90.1 <math>\pm</math> 3.4</u>  | <b>90.5 <math>\pm</math> 3.2</b>  | 90.0 $\pm$ 3.5                    |
| GST [PDw]            | 38.1 $\pm$ 16.6                   | <u>46.4 <math>\pm</math> 15.9</u> | <b>46.6 <math>\pm</math> 12.2</b> | <b>74.9 <math>\pm</math> 9.0</b>  | <u>70.8 <math>\pm</math> 12.8</u> | 69.4 $\pm$ 11.2                   | <b>90.1 <math>\pm</math> 4.7</b>  | 89.5 $\pm$ 5.1                    | <u>90.1 <math>\pm</math> 4.7</u>  |
| <b>Out of Domain</b> |                                   |                                   |                                   |                                   |                                   |                                   |                                   |                                   |                                   |
| HH [T1w]             | <u>49.4 <math>\pm</math> 14.1</u> | 33.0 $\pm$ 14.2                   | <b>57.7 <math>\pm</math> 13.8</b> | <b>63.0 <math>\pm</math> 11.7</b> | 59.3 $\pm$ 12.0                   | <u>61.1 <math>\pm</math> 13.2</u> | <b>81.6 <math>\pm</math> 6.5</b>  | 75.5 $\pm$ 10.9                   | <u>77.4 <math>\pm</math> 8.9</u>  |
| HH [T2w]             | <b>58.6 <math>\pm</math> 7.9</b>  | <u>53.8 <math>\pm</math> 6.7</u>  | 46.5 $\pm$ 7.0                    | <b>75.0 <math>\pm</math> 9.1</b>  | <u>72.0 <math>\pm</math> 8.1</u>  | 65.6 $\pm$ 7.9                    | 87.2 $\pm$ 6.4                    | <b>89.7 <math>\pm</math> 5.8</b>  | <u>88.1 <math>\pm</math> 6.1</u>  |
| HH [PDw]             | 33.8 $\pm$ 19.3                   | <u>39.4 <math>\pm</math> 16.2</u> | <b>40.3 <math>\pm</math> 15.0</b> | <u>60.5 <math>\pm</math> 15.0</u> | <b>61.5 <math>\pm</math> 13.4</b> | 59.7 $\pm$ 16.4                   | 82.7 $\pm$ 10.0                   | <u>83.1 <math>\pm</math> 10.5</u> | <b>85.6 <math>\pm</math> 8.5</b>  |
| IOP [T1w]            | <u>50.6 <math>\pm</math> 19.5</u> | 30.7 $\pm$ 17.9                   | <b>54.4 <math>\pm</math> 15.5</b> | <u>58.3 <math>\pm</math> 15.8</u> | <b>60.9 <math>\pm</math> 18.0</b> | 57.4 $\pm$ 20.1                   | <b>79.1 <math>\pm</math> 13.2</b> | 70.7 $\pm$ 16.4                   | <u>74.0 <math>\pm</math> 13.5</u> |
| IOP [T2w]            | <b>58.3 <math>\pm</math> 9.4</b>  | <u>43.8 <math>\pm</math> 8.3</u>  | 40.6 $\pm$ 12.8                   | <b>74.7 <math>\pm</math> 6.1</b>  | <u>71.4 <math>\pm</math> 5.9</u>  | 63.6 $\pm$ 10.0                   | 85.1 $\pm$ 4.8                    | <u>85.8 <math>\pm</math> 5.1</u>  | <b>86.1 <math>\pm</math> 4.4</b>  |
| IOP [PDw]            | 31.1 $\pm$ 20.8                   | <u>36.6 <math>\pm</math> 21.3</u> | <b>37.6 <math>\pm</math> 16.4</b> | <u>59.2 <math>\pm</math> 13.5</u> | 55.3 $\pm$ 16.2                   | <b>59.7 <math>\pm</math> 14.1</b> | <u>76.4 <math>\pm</math> 10.5</u> | 76.3 $\pm$ 11.5                   | <b>77.2 <math>\pm</math> 9.9</b>  |

### 4.3.2. STROKE LESION SEGMENTATION

We next evaluate on the ARC dataset (Gibson et al., 2024) using both DSC and HD95 (see Table 2). **SeqInv** achieves the best overall performance on T1w, improving DSC by 0.5 points and reducing HD95 by 5.9 mm compared to the baseline. On T2w, **SeqAug** reduces HD95 by 22.2 mm, indicating excellent boundary accuracy while maintaining a competitive DSC. For FLAIR, **SeqInv** provides a further 4.7 mm decrease in HD95, offering improved boundary delineation over the baseline.

Table 2: Stroke lesion segmentation performance using 100% training data. Values show mean  $\pm$  std, with **bold** and underlined indicating best and second-best results for each metric. DSC (higher is better) and HD95 in mm (lower is better) are shown for each model.

|             | DSC ( $\uparrow$ )                |                                   |                                   | HD95 (mm) ( $\downarrow$ )        |                                   |                                   |
|-------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
|             | Base                              | SeqAug                            | SeqInv                            | Base                              | SeqAug                            | SeqInv                            |
| ARC [T1w]   | <u>78.4 <math>\pm</math> 12.8</u> | 77.3 $\pm$ 14.6                   | <b>78.9 <math>\pm</math> 12.2</b> | <u>33.2 <math>\pm</math> 26.4</u> | 36.3 $\pm$ 31.0                   | <b>27.3 <math>\pm</math> 23.7</b> |
| ARC [T2w]   | 78.7 $\pm$ 11.2                   | <b>80.3 <math>\pm</math> 9.4</b>  | <u>79.4 <math>\pm</math> 11.1</u> | 36.2 $\pm$ 26.5                   | <b>14.0 <math>\pm</math> 12.7</b> | <u>24.5 <math>\pm</math> 24.8</u> |
| ARC [FLAIR] | 68.4 $\pm$ 26.7                   | <u>71.0 <math>\pm</math> 22.5</u> | <b>71.1 <math>\pm</math> 23.0</b> | <u>67.9 <math>\pm</math> 20.5</u> | 68.1 $\pm$ 18.4                   | <b>63.2 <math>\pm</math> 14.0</b> |

### 4.3.3. MRI DENOISING

Lastly, we evaluate PSNR on IXI volumes corrupted with synthetic noise (Table 3). **SeqInv** achieves notable gains on T1w, boosting PSNR by up to 4.2 dB with only 1% training data, and these gains persist even at 100% training data, suggesting robust feature learning. Out-of-domain generalisation is also particularly strong, with **SeqInv** reaching 21.7 dB on HH T1w compared to 19.3 dB for the baseline. By contrast, **SeqAug** provides moderate gains, indicating that purely contrast-based augmentation alone cannot match the full sequence-invariant approach.

Table 3: Image denoising performance using Peak Signal-to-Noise Ratio in dB (higher is better). Values show mean  $\pm$  std, with **bold** and underlined indicating best and second-best results. GST represents the training domain.

|                      | 1% Training Data                 |                                  |                                  | 10% Training Data                |                                  |                                  | 100% Training Data |                                  |                                  |
|----------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|--------------------|----------------------------------|----------------------------------|
|                      | Base                             | SeqAug                           | SeqInv                           | Base                             | SeqAug                           | SeqInv                           | Base               | SeqAug                           | SeqInv                           |
| <b>In Domain</b>     |                                  |                                  |                                  |                                  |                                  |                                  |                    |                                  |                                  |
| GST [T1w]            | 14.9 $\pm$ 0.2                   | <u>16.2 <math>\pm</math> 0.3</u> | <b>19.1 <math>\pm</math> 0.4</b> | 19.0 $\pm$ 0.4                   | <u>19.7 <math>\pm</math> 0.8</u> | <b>20.3 <math>\pm</math> 1.0</b> | 19.1 $\pm$ 0.3     | <u>20.6 <math>\pm</math> 0.4</u> | <b>21.3 <math>\pm</math> 0.4</b> |
| GST [T2w]            | 17.2 $\pm$ 0.2                   | <b>17.7 <math>\pm</math> 0.2</b> | <u>17.3 <math>\pm</math> 0.3</u> | 18.3 $\pm$ 0.3                   | <u>18.5 <math>\pm</math> 0.6</u> | <b>19.8 <math>\pm</math> 0.3</b> | 18.3 $\pm$ 0.3     | <u>19.4 <math>\pm</math> 0.3</u> | <b>20.0 <math>\pm</math> 0.4</b> |
| GST [PDw]            | 17.0 $\pm$ 0.2                   | <u>18.3 <math>\pm</math> 0.3</u> | <b>18.7 <math>\pm</math> 0.3</b> | 18.6 $\pm$ 0.3                   | <u>19.3 <math>\pm</math> 0.5</u> | <b>20.0 <math>\pm</math> 0.5</b> | 18.7 $\pm$ 0.3     | <u>19.9 <math>\pm</math> 0.3</u> | <b>20.6 <math>\pm</math> 0.3</b> |
| <b>Out of Domain</b> |                                  |                                  |                                  |                                  |                                  |                                  |                    |                                  |                                  |
| HH [T1w]             | 15.1 $\pm$ 0.1                   | <u>16.5 <math>\pm</math> 0.2</u> | <b>19.4 <math>\pm</math> 0.2</b> | 19.1 $\pm$ 0.3                   | <u>20.0 <math>\pm</math> 0.7</u> | <b>20.1 <math>\pm</math> 1.1</b> | 19.3 $\pm$ 0.2     | <u>21.0 <math>\pm</math> 0.3</u> | <b>21.7 <math>\pm</math> 0.3</b> |
| HH [T2w]             | 16.5 $\pm$ 0.4                   | <b>16.9 <math>\pm</math> 0.5</b> | 16.4 $\pm$ 0.5                   | <u>17.5 <math>\pm</math> 0.4</u> | 15.6 $\pm$ 1.4                   | <b>18.8 <math>\pm</math> 0.6</b> | 17.5 $\pm$ 0.4     | <u>18.5 <math>\pm</math> 0.6</u> | <b>18.9 <math>\pm</math> 0.6</b> |
| HH [PDw]             | 16.5 $\pm$ 0.4                   | <b>17.8 <math>\pm</math> 0.5</b> | <u>17.8 <math>\pm</math> 0.6</u> | 18.0 $\pm$ 0.5                   | <b>19.2 <math>\pm</math> 0.6</b> | <u>18.9 <math>\pm</math> 1.0</u> | 18.2 $\pm$ 0.5     | <u>19.3 <math>\pm</math> 0.6</u> | <b>19.9 <math>\pm</math> 0.7</b> |
| IOP [T1w]            | 14.7 $\pm$ 0.2                   | <u>16.7 <math>\pm</math> 0.2</u> | <b>18.9 <math>\pm</math> 0.2</b> | 18.4 $\pm$ 0.2                   | <b>19.9 <math>\pm</math> 0.3</b> | <u>18.5 <math>\pm</math> 0.9</u> | 18.8 $\pm$ 0.1     | <u>20.3 <math>\pm</math> 0.2</u> | <b>21.0 <math>\pm</math> 0.2</b> |
| IOP [T2w]            | <u>17.1 <math>\pm</math> 0.1</u> | <b>17.6 <math>\pm</math> 0.1</b> | 17.0 $\pm$ 0.1                   | 17.9 $\pm$ 0.2                   | <u>18.8 <math>\pm</math> 0.3</u> | <b>19.6 <math>\pm</math> 0.2</b> | 18.0 $\pm$ 0.1     | <u>19.2 <math>\pm</math> 0.1</u> | <b>19.7 <math>\pm</math> 0.2</b> |
| IOP [PDw]            | 16.9 $\pm$ 0.1                   | <u>18.3 <math>\pm</math> 0.1</u> | <b>18.8 <math>\pm</math> 0.1</b> | 18.5 $\pm$ 0.1                   | <u>19.8 <math>\pm</math> 0.2</u> | <b>20.0 <math>\pm</math> 0.3</b> | 18.6 $\pm$ 0.1     | <u>19.8 <math>\pm</math> 0.1</u> | <b>20.5 <math>\pm</math> 0.2</b> |

## 5. Discussion

Our results show that sequence-invariant self-supervised learning substantially improves model robustness and generalisation across diverse MRI sequences and acquisition sites. In particular, it enables effective feature learning even with minimal labelled data, suggesting that the method captures fundamental anatomical cues independent of sequence-specific contrast.

### 5.1. Key Findings

We highlight three key aspects of our method’s performance. First, even when trained on as little as 1% of the data, it achieves up to +4.2 dB PSNR in denoising and +8.3 DSC points in segmentation, underscoring its robust representation capabilities. Second, the model generalises well across T1w, T2w, and PDw, showing particularly strong results on T1w while leaving some gaps on the other sequences. Finally, it excels at out-of-domain adaptation, often surpassing baseline models more in unseen sites than in the original training domain, illustrating its effectiveness for cross-site generalisation.

### 5.2. Limitations

Our approach faces several limitations. First, training on full-resolution 3D volumes is computationally demanding, limiting batch size and reducing the number of negative pairs used during contrastive learning. Second, we rely on a CNN backbone, which may not capture long-range dependencies as effectively as vision transformers or other recent architectures. Third, while cross-sequence invariance bolsters model robustness, certain sequence-specific gaps - particularly on T2w images - highlight the need for further improvements. Finally, qMRI data are relatively scarce, leading us to pre-train on only 51 subjects; scaling to a larger qMRI dataset or inferring parameter maps from a resource like the UK Biobank (Sudlow et al., 2015) could significantly boost downstream performance.

### 5.3. Future Directions

We plan to explore ViT-based encoders, alternative SSL methods (VICReg, DINO), and larger clinical datasets to further improve domain invariance and push towards routine clinical adoption. We also expect multi-view contrastive learning and decoder pretraining (Asiedu et al., 2022) to be valuable directions of future work.

### 5.4. Conclusion

Sequence-invariant self-supervised learning offers a promising route towards more robust, generalisable medical image analysis. By using physics-informed contrast simulation and contrastive training, we can exploit the shared anatomy across varied MRI sequences and sites. Although challenges remain - especially around computational cost and data availability - our results illustrate the potential for significant gains in low-data scenarios and out-of-domain adaptation. We believe this framework provides a stepping stone toward truly cross-domain, clinically deployable deep learning models in medical imaging.



## Acknowledgments

LC is supported by the EPSRC-funded UCL Centre for Doctoral Training in Intelligent, Integrated Imaging in Healthcare (i4health) (EP/S021930/1), and the Wellcome Trust (203147/Z/16/Z and 205103/Z/16/Z). This research was supported by NVIDIA and used NVIDIA RTX A6000 48GB.

## References

- Rhea Adams, Walter Zhao, Siyuan Hu, Wenjiao Lyu, Khoi Minh Huynh, Sahar Ahmad, Dan Ma, and Pew-Thian Yap. Ultimatesynth: Mri physics for pan-contrast ai. December 2024. doi: 10.1101/2024.12.05.627056. URL <http://dx.doi.org/10.1101/2024.12.05.627056>.
- Emmanuel Brempong Asiedu, Simon Kornblith, Ting Chen, Niki Parmar, Matthias Minderer, and Mohammad Norouzi. Decoder denoising pretraining for semantic segmentation, 2022. URL <https://arxiv.org/abs/2205.11423>.
- Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John Guttag, and Adrian V. Dalca. Voxelmorph: A learning framework for deformable medical image registration. *IEEE Transactions on Medical Imaging*, 38(8):1788–1800, August 2019. ISSN 1558-254X. doi: 10.1109/tmi.2019.2897538. URL <http://dx.doi.org/10.1109/TMI.2019.2897538>.
- Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for self-supervised learning, 2022. URL <https://arxiv.org/abs/2105.04906>.
- Benjamin Billot, Douglas N. Greve, Oula Puonti, et al. SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical Image Analysis*, 86:102789, May 2023. ISSN 1361-8415. doi: 10.1016/j.media.2023.102789.
- Pedro Borges, Richard Shaw, Thomas Varsavsky, et al. Acquisition-invariant brain mri segmentation with informative uncertainties, 2021.
- Krishna Chaitanya, Ertunc Erdil, Neerav Karani, and Ender Konukoglu. Contrastive learning of global and local features for medical image segmentation with limited annotations, 2020. URL <https://arxiv.org/abs/2006.10511>.
- Liam Chalcraft, Ruben Lourenço Pereira, Mikael Brudfors, Andrew S. Kayser, Mark D’Esposito, Cathy J. Price, Ioannis Pappas, and John Ashburner. Large-kernel attention for efficient and robust brain lesion segmentation, 2023. URL <https://arxiv.org/abs/2308.07251>.
- Liam Chalcraft, Jenny Crinion, Cathy J. Price, and John Ashburner. Domain-agnostic stroke lesion segmentation using physics-constrained synthetic data, 2024. URL <https://arxiv.org/abs/2412.03318>.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations, 2020. URL <https://arxiv.org/abs/2002.05709>.

- Reuben Dorent, Aaron Kujawa, Marina Ivory, et al. Crossmoda 2021 challenge: Benchmark of cross-modality domain adaptation techniques for vestibular schwannoma and cochlea segmentation. *Medical Image Analysis*, 83:102628, January 2023. ISSN 1361-8415. doi: 10.1016/j.media.2022.102628. URL <http://dx.doi.org/10.1016/j.media.2022.102628>.
- Makayla Gibson, Roger Newman-Norlund, Leonardo Bonilha, Julius Fridriksson, Gregory Hickok, Argye E. Hillis, Dirk Bart den Ouden, and Christopher Rorden. The Aphasia Recovery Cohort, an open-source chronic stroke repository. *Scientific Data*, 11(1):1–8, 2024. ISSN 20524463. doi: 10.1038/s41597-024-03819-7. URL <http://dx.doi.org/10.1038/s41597-024-03819-7>.
- Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning, 2020. URL <https://arxiv.org/abs/2006.07733>.
- HáKon Gudbjartsson and Samuel Patz. The rician distribution of noisy mri data. *Magnetic Resonance in Medicine*, 34(6):910–914, December 1995. ISSN 1522-2594. doi: 10.1002/mrm.1910340618. URL <http://dx.doi.org/10.1002/mrm.1910340618>.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning, 2020. URL <https://arxiv.org/abs/1911.05722>.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners, 2021. URL <https://arxiv.org/abs/2111.06377>.
- Shih-Cheng Huang, Anuj Pareek, Malte Jensen, Matthew P. Lungren, Serena Yeung, and Akshay S. Chaudhari. Self-supervised learning for medical image classification: a systematic review and implementation guidelines. *npj Digital Medicine*, 6(1), April 2023. ISSN 2398-6352. doi: 10.1038/s41746-023-00811-0. URL <http://dx.doi.org/10.1038/s41746-023-00811-0>.
- Olivier J. Hénaff, Aravind Srinivas, Jeffrey De Fauw, Ali Razavi, Carl Doersch, S. M. Ali Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding, 2020. URL <https://arxiv.org/abs/1905.09272>.
- Wenxuan Li, Chongyu Qu, Xiaoxi Chen, Pedro R. A. S. Bassi, Yijia Shi, Yuxiang Lai, Qian Yu, Huimin Xue, Yixiong Chen, Xiaorui Lin, Yutong Tang, Yining Cao, Haoqi Han, Zheyuan Zhang, Jiawei Liu, Tiezheng Zhang, Yujiu Ma, Jincheng Wang, Guang Zhang, Alan Yuille, and Zongwei Zhou. Abdomenatlas: A large-scale, detailed-annotated, & multi-center dataset for efficient transfer learning and open algorithmic benchmarking, 2024. URL <https://arxiv.org/abs/2407.16697>.
- Peirong Liu, Oula Puonti, Xiaoling Hu, Daniel C. Alexander, and Juan E. Iglesias. Brainid: Learning contrast-agnostic anatomical representations for brain imaging, 2024a. URL <https://arxiv.org/abs/2311.16914>.

- Peirong Liu, Oula Puonti, Annabel Sorby-Adams, William T. Kimberly, and Juan E. Iglesias. Pepsi: Pathology-enhanced pulse-sequence-invariant representations for brain mri, 2024b. URL <https://arxiv.org/abs/2403.06227>.
- Dan Ma, Vikas Gulani, Nicole Seiberlich, Kecheng Liu, Jeffrey L. Sunshine, Jeffrey L. Duerk, and Mark A. Griswold. Magnetic resonance fingerprinting. *Nature*, 495(7440):187–192, March 2013. ISSN 1476-4687. doi: 10.1038/nature11971. URL <http://dx.doi.org/10.1038/nature11971>.
- Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15(1), January 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-44824-z. URL <http://dx.doi.org/10.1038/s41467-024-44824-z>.
- Maria Ines Meyer, Ezequiel de la Rosa, Nuno Pedrosa de Barros, Roberto Paoella, Koen Van Leemput, and Diana M. Sima. A contrast augmentation approach to improve multi-scanner generalization in mri. *Frontiers in Neuroscience*, 15, August 2021. ISSN 1662-453X. doi: 10.3389/fnins.2021.708196. URL <http://dx.doi.org/10.3389/fnins.2021.708196>.
- John N. Morelli, Val M. Runge, Fei Ai, Ulrike Attenberger, Lan Vu, Stuart H. Schmeets, Wolfgang R. Nitz, and John E. Kirsch. An image-based approach to understanding the physics of mr artifacts. *RadioGraphics*, 31(3):849–866, May 2011. ISSN 1527-1323. doi: 10.1148/rg.313105115. URL <http://dx.doi.org/10.1148/rg.313105115>.
- Abhishek Narayanan, Rushabh Musthyala, Rahul Sankar, Anirudh Prasad Nistala, Pranav Singh, and Jacopo Cirrone. Free form medical visual question answering in radiology, 2024. URL <https://arxiv.org/abs/2401.13081>.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khali-dov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2024. URL <https://arxiv.org/abs/2304.07193>.
- Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting, 2016. URL <https://arxiv.org/abs/1604.07379>.
- Emma C. Robinson, Alexander Hammers, Anders Ericsson, A. David Edwards, and Daniel Rueckert. Identifying population differences in whole-brain structural networks: A machine learning approach. *NeuroImage*, 50(3):910–919, April 2010. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2010.01.019. URL <http://dx.doi.org/10.1016/j.neuroimage.2010.01.019>.
- Mélanie Roschewitz, Fabio De Sousa Ribeiro, Tian Xia, Galvin Khara, and Ben Glocker. Robust image representations with counterfactual contrastive learning, 2024. URL <https://arxiv.org/abs/2409.10365>.

- Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, Bette Liu, Paul Matthews, Giok Ong, Jill Pell, Alan Silman, Alan Young, Tim Sprosen, Tim Peakman, and Rory Collins. Uk biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLOS Medicine*, 12(3):e1001779, March 2015. ISSN 1549-1676. doi: 10.1371/journal.pmed.1001779. URL <http://dx.doi.org/10.1371/journal.pmed.1001779>.
- L.N. Tanenbaum, A.J. Tsiouris, A.N. Johnson, T.P. Naidich, M.C. DeLano, E.R. Melhem, P. Quarterman, S.X. Parameswaran, A. Shankaranarayanan, M. Goyen, and A.S. Field. Synthetic mri for clinical neuroimaging: Results of the magnetic resonance image compilation (magic) prospective, multicenter, multireader trial. *American Journal of Neuroradiology*, 38(6):1103–1110, April 2017. ISSN 1936-959X. doi: 10.3174/ajnr.a5227. URL <http://dx.doi.org/10.3174/ajnr.A5227>.
- Yucheng Tang, Dong Yang, Wenqi Li, Holger Roth, Bennett Landman, Daguang Xu, Vishwesh Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical image analysis, 2022. URL <https://arxiv.org/abs/2111.14791>.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding, 2019. URL <https://arxiv.org/abs/1807.03748>.
- K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens. Automated model-based tissue classification of mr images of the brain. *IEEE Transactions on Medical Imaging*, 18(10):897–908, 1999. ISSN 0278-0062. doi: 10.1109/42.811270. URL <http://dx.doi.org/10.1109/42.811270>.
- Dequan Wang, Evan Shelhamer, Shaoteng Liu, et al. Tent: Fully test-time adaptation by entropy minimization, 2021.
- Nikolaus Weiskopf, Luke J. Edwards, Gunther Helms, Siawoosh Mohammadi, and Evgeniya Kirilina. Quantitative magnetic resonance imaging of brain anatomy and in vivo histology. *Nature Reviews Physics*, 3(8):570–588, June 2021. ISSN 2522-5820. doi: 10.1038/s42254-021-00326-1. URL <http://dx.doi.org/10.1038/s42254-021-00326-1>.
- Linshan Wu, Jiaxin Zhuang, and Hao Chen. Large-scale 3d medical image pre-training with geometric context priors, 2024. URL <https://arxiv.org/abs/2410.09890>.
- Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling, 2022. URL <https://arxiv.org/abs/2111.09886>.
- Yue Yang, Mona Gandhi, Yufei Wang, Yifan Wu, Michael S. Yao, Chris Callison-Burch, James C. Gee, and Mark Yatskar. A textbook remedy for domain shifts: Knowledge priors for medical image analysis, 2024. URL <https://arxiv.org/abs/2405.14839>.
- Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction, 2021. URL <https://arxiv.org/abs/2103.03230>.

## Appendix A. Physics-based Signal Equations

This appendix details the forward models used to simulate MRI signal intensities from quantitative MRI (qMRI) parameter maps. For each voxel, we assume known proton density (PD), longitudinal relaxation rate ( $R_1$ ), and transverse relaxation rate ( $R_2$  or  $R_2^*$ ), and optionally a magnetic transfer (MT) parameter. We consider several common MRI sequences, each governed by its own signal equation:

### Fast Spin-Echo (FSE).

$$S_{\text{FSE}} = B_1 \cdot \text{PD} \cdot (1 - e^{-R_1 \cdot T_R}) \cdot e^{-R_2 \cdot T_E}. \quad (1)$$

### Gradient-Echo (GRE).

$$S_{\text{GRE}} = B_1 \cdot \text{PD} \cdot \sin(\alpha) \cdot (1 - \text{MT}) \cdot \frac{1 - e^{-R_1 \cdot T_R}}{1 - \cos(\alpha)(1 - \text{MT}) e^{-R_1 \cdot T_R}} \cdot e^{-R_2^* \cdot T_E}. \quad (2)$$

### Fluid-Attenuated Inversion Recovery (FLAIR).

$$S_{\text{FLAIR}} = B_1 \cdot \text{PD} \cdot e^{-R_2 \cdot T_E} \cdot (1 - 2e^{-R_1 \cdot T_I} + e^{-R_1 \cdot T_R}). \quad (3)$$

### Magnetisation-Prepared Rapid Gradient Echo (MPRAGE).

$$S_{\text{MPRAGE}} = \left| \sin(\alpha) \left[ \frac{1 - e^{-R_1 \cdot T_R}}{1 - \cos(\alpha) e^{-R_1 \cdot T_R}} (1 - (\cos(\alpha) e^{-T_X \cdot R_1})^n) \right] e^{-T_D \cdot R_1} + (1 - e^{-T_D \cdot R_1}) \right|. \quad (4)$$

Here,  $B_1$  represents the receive field strength,  $\alpha$  (in radians) is the flip angle,  $T_R$  the repetition time,  $T_E$  the echo time,  $T_I$  the inversion time,  $T_X$  the excitation repetition time,  $T_D$  a delay time, and  $n$  the number of excitation pulses (all times in seconds unless noted otherwise).

### Noise Simulation

To introduce realistic MRI noise, we add Rician noise (Gudbjartsson and Patz, 1995) to the simulated signals. Specifically,

$$S_{\text{noisy}} = \sqrt{(S_{\text{MRI}} + \epsilon_r)^2 + \epsilon_i^2} \text{ where } \epsilon_r \sim \mathcal{N}(0, \sigma^2) \text{ and } \epsilon_i \sim \mathcal{N}(0, \sigma^2) \quad (5)$$

where  $S_{\text{MRI}}$  is the noise-free signal, and each  $\mathcal{N}(0, \sigma^2)$  term is a zero-mean Gaussian simulating real and imaginary noise components.

All simulations were implemented using the NiTorch library<sup>†</sup>.

### Acquisition Parameters and Data Augmentation

We randomly sampled MRI sequences and parameters from realistic distributions to produce diverse training data, as summarised below. Each sequence was chosen with equal probability to ensure broad coverage of acquisition protocols.

<sup>†</sup><https://github.com/balbasty/nitorch>

**FLAIR:**

$$T_E \sim 10^{\mathcal{N}(\log_{10}(0.02), \log_{10}(0.1))} \text{ s}, \quad T_R \sim 10^{\mathcal{U}(\log_{10}(0.001), \log_{10}(5))} \text{ s},$$

$$T_I \sim 10^{\mathcal{U}(\log_{10}(0.001), \log_{10}(3))} \text{ s}.$$

**FSE:**

$$T_E \sim 10^{\mathcal{U}(\log_{10}(0.001), \log_{10}(3))} \text{ s}, \quad T_R \sim 10^{\mathcal{U}(\log_{10}(0.001), \log_{10}(3))} \text{ s}.$$

**MPRAGE:**

$$T_R \sim \mathcal{N}(23, 2.3) \text{ s}, \quad T_I \sim \mathcal{U}(0.6, 0.9) \text{ s}, \quad T_X \sim \mathcal{U}(0.004, 0.008) \text{ s},$$

$$T_E \sim \mathcal{U}(0.002, 0.004) \text{ s}, \quad \alpha \sim \mathcal{U}(5^\circ, 12^\circ).$$

**GRE:**

$$T_E \sim 10^{\mathcal{U}(\log_{10}(0.002), \log_{10}(0.08))} \text{ s}, \quad T_R \sim 10^{\mathcal{U}(\log_{10}(0.005), \log_{10}(5))} \text{ s}, \quad \alpha \sim \mathcal{U}(5^\circ, 50^\circ).$$

All sampled parameters were constrained to physically plausible ranges (e.g. taking absolute values as needed). This broad variation in sequences and parameter values helps train the model to handle diverse imaging protocols.

## Appendix B. Model Architectures

### Pre-training Architecture

The pre-training setup consists of three components:

- **CNN Encoder:**
  - Input: 3D volume with 1 channel
  - 5 downsampling blocks with features: (64, 128, 256, 512, 768)
  - Each block: Two 3D convolutions with kernel size 3
  - Instance normalization and GELU activation after each convolution
  - Dropout rate: 0.2
- **Projector:** Two-layer MLP
  - Input dimension: 768
  - Hidden layer (768  $\rightarrow$  512) with GELU activation
  - Output projects to contrastive learning space (512  $\rightarrow$  128)
- **Reconstructor:** Decoder network
  - 4 transposed convolution blocks: (768  $\rightarrow$  384  $\rightarrow$  192  $\rightarrow$  96  $\rightarrow$  48)
  - Each upsampling:  $2 \times 2 \times 2$  stride
  - Final  $1 \times 1 \times 1$  convolution to output channels

## Downstream Task Architectures

For the denoising task, we use a U-Net architecture that incorporates the pre-trained encoder:

- **CNN U-Net:**
  - Input: 3D volume with 1 channel
  - Encoder: Pre-trained CNN encoder (frozen)
  - Feature dimensions: (768, 512, 256, 128, 64, 32)
  - Instance normalization throughout
  - GELU activation functions
  - Dropout rate: 0.2
  - Upsampling: Transposed convolutions
  - Output: 1 channel (predicted noise)

### B.1. Training Details

The models were trained with the following specifications:

- Optimizer: AdamW with gradient clipping at 12.0
- Learning rate schedule:  $(1 - \frac{epoch}{max\_epochs})^{0.9}$
- Loss functions:
  - Pre-training: NT-Xent loss + L1 reconstruction loss
  - Denoising: Mean Squared Error (MSE)
  - Segmentation: Dice + Cross-Entropy
- Patch size:  $96 \times 96 \times 96$
- Mixed precision training
- Batch size:
  - Pre-training: 8
  - Downstream tasks: 2

During downstream task training, the pre-trained encoder weights were frozen while the decoder weights were trained from scratch, as evidenced by the weight loading and gradient freezing in the training code.