# Self-Supervised Learning on 3D Brain MRI Data

Angelo Jovin Yamachui Sitcheu  Moussa Kassem Sbeyti
Ralf Mikut  Nadja Klein

Start on April 2026

3D Magnetic Resonance Imaging (MRI) is central to modern biomedical research and clinical workflows. However, developing high-performing 3D medical vision models typically requires large amounts of annotated data, especially for dense prediction tasks such as organ or structure segmentation. Obtaining voxel-wise annotations at scale is expensive, slow, and often infeasible.

Self-supervised learning (SSL) has emerged as a promising approach to leverage large collections of unlabeled medical volumes by learning general-purpose representations that can be adapted to downstream tasks with limited labels. The OpenMind dataset (see Figure 1) provides a large-scale, diverse corpus of 3D MRI volumes collected across many scanners and MRI modalities, enabling systematic investigation of scalable SSL for 3D medical imaging in a standardized setting.

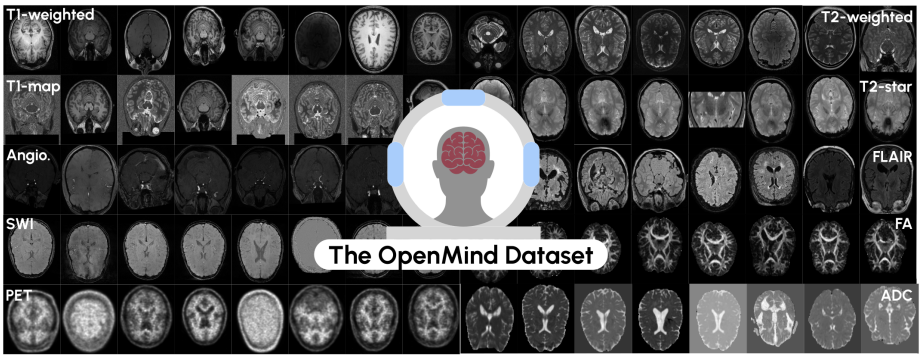


Figure 1: Example MRI samples from the OpenMind dataset (illustrative).

This thesis aims to investigate SSL pre-training strategies for 3D MRI on the OpenMind dataset, with a particular focus on predictive SSL objectives inspired by JEPA (Joint Embedding Predictive Architectures) and on efficient non-standard backbone families such as state-space / Mamba-style vision models. Following the evaluation protocols, the thesis will emphasize SSL methods for competitive downstream analyses in segmentation as well as classification.

The thesis will evaluate existing SSL baselines and explore design choices specific to 3D medical volumes, resulting in a reproducible benchmark and a prototype method or training recipe.

Objectives:

- Conduct a literature review on SSL for 3D medical imaging, including masked modeling, contrastive methods, and predictive/JEPA-style approaches.
- Study architectural alternatives to CNNs/ViTs for volumetric representation learning, with emphasis on efficient backbones (e.g., state-space / Vision Mamba-style models) and how they can be adapted to 3D data.
- Establish an experimental baseline on OpenMind by reproducing and/or re-using existing SSL pre-training setups and evaluating downstream performance.
- Evaluate representations on downstream tasks segmentation and classification, and analyze robustness across dataset.

Requirements:

- Currently enrolled in a Master's program in Computer Science (Informatik) at KIT, or in a related engineering field.
- Solid knowledge of deep learning for computer vision and image processing
- Advanced programming experience in Python
- Familiarity with relevant libraries (e.g., PyTorch, MONAI, NumPy)
- (Nice-to-have) Experience with medical imaging formats and tooling (e.g., NIfTI, SimpleITK)

Benefits:

- The thesis can be written in English or German
- Possibility to work remotely
- Structured onboarding and continuous support
- Contribution to an active research area: scalable and robust SSL for 3D medical imaging